ADVANCED IPv6 ENGINEERING

Design, Deploy, Secure, and Operate IPv6 Networks at Scale

DESIGN | DEPLOY | SECURE | OPERATE

CHRISTOPHER HENDERSON

Advanced IPv6 Engineering

Design, Deploy, Secure, and Operate IPv6 Networks at Scale

Steve Publications

This book is available at https://leanpub.com/advancedipv6engineering

This version was published on 2026-09-07

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

Contents

Design, Deploy, Secure, and Operate IPv6 Networks at Scale

This book is a comprehensive, technically rigorous guide to IPv6 for experienced network engineers, systems engineers, and infrastructure architects. It covers the full lifecycle of IPv6: from addressing design and routing protocol configuration to security architecture, operational observability, automation, and large-scale migration from IPv4. Each chapter presents the underlying protocol mechanics, architectural tradeoffs, and production-ready configuration examples for major platforms including Linux, enterprise network equipment, cloud providers, containers, and Kubernetes. The material assumes familiarity with fundamental networking concepts and is written to serve as both a learning resource and a long-term reference for real-world deployments.

Introduction

The IPv4 Exhaustion That Changed Everything

The Internet has always borrowed time. When the original designers allocated 32-bit addresses, they anticipated a few thousand networks and a few million hosts. They got 4.3 billion addresses and far more demand. Classful addressing squandered blocks of 16 million addresses on networks needing a dozen. The 1990s panic that followed birthed CIDR, NAT, and carrier-grade NAT. Each extension bought another decade. By 2011 the IANA public pool ran dry. Regional Internet Registries followed in rapid succession. The era of freely available IPv4 addresses ended.

What emerged from that exhaustion was a fundamental realization: NAT is a hack that works well enough for connectivity but breaks end-to-end principles, adds operational complexity, and creates security false promises. IPv6 is not merely a larger address space. It is a protocol redesign that addresses the limitations baked into IPv4 from the start, including header inefficiency, lack of mandatory security integration, fragile extension mechanisms, and a Neighbor Discovery architecture that never existed because ARP was bolted on decades later.

The transition has been agonizingly slow. That slowness is neither accidental nor inevitable. It is the product of technical friction, organizational inertia, and a dual-stack coexistence model that allows many organizations to delay indefinitely. But the clock is ticking. Mobile networks increasingly deploy IPv6-first or IPv6-only access. Regulatory mandates in countries including India, China, and parts of the European Union have accelerated adoption. Cloud providers require IPv6-ready infrastructure. The question is no longer whether you need IPv6. The question is whether you are ready when your customers, regulators, or peers stop accommodating IPv4-only.

Where IPv6 Stands Today

Global IPv6 adoption crossed a meaningful threshold in 2026. Google statistics show global IPv6 traffic at or above 50 percent for the first time, up from

roughly 43 percent in early 2025 [1]. APNIC economies have maintained IPv6 capability above 50 percent since April 2025, representing more than 64 percent of global IPv6 users [2]. France leads national adoption at roughly 85 percent on Google traffic. Germany, India, and the United States have each surpassed 50 percent [3].

These numbers reflect only one vantage point. The reality is more nuanced. Mobile networks in many regions have reached 90 percent plus IPv6 capability. Residential broadband is catching up. Enterprise internal networks lag significantly. Large organizations that started planning five years ago are now in advanced deployment phases. Many others are still in the education stage, running IPv6 only on a few external web servers.

The state of the Internet as of 2026 is best described as bi-modal. A core set of technology-forward organizations and mobile operators are operating mature, large-scale IPv6 environments. A larger set of enterprises is managing fragile dual-stack networks that work most of the time but fail in unpredictable ways under certain conditions. And a substantial number of organizations have IPv6 partially enabled on some systems, creating connectivity gaps and security blind spots that they are unaware of.

This book is written for engineers who want to be in the first category and avoid the pitfalls of the second and third.

What This Book Covers

This is an advanced book. It assumes you already understand Ethernet, TCP/IP, routing basics, DNS fundamentals, and firewall concepts. You should be comfortable reading configuration files, running diagnostic commands, and thinking in terms of network architecture. If you are a CCNP-level network engineer, a senior Linux systems administrator, or a cloud infrastructure architect, this book should meet you where you are.

The material is organized into several parts that build progressively but are also designed as a reference. Early chapters cover IPv6 fundamentals from an advanced perspective: addressing structure, subnetting design, address planning and allocation, and the mechanisms that replace IPv4 constructs. Middle chapters cover routing protocols, dual-stack and IPv6-only architectures, transition mechanisms, and translation technologies. Later chapters address security, monitoring, automation, and migration strategy.

Throughout the book you will find:

- Detailed protocol explanations that show not just what happens but why it happens that way
- RFC references for every protocol behavior, clearly distinguished from implementation-specific details
- Complete, copy-paste-ready configuration examples for Linux, enterprise routers, firewalls, cloud platforms, containers, and Kubernetes
- Verification procedures that show how to confirm configurations are working as intended
- Real-world deployment patterns and known pitfalls
- Security considerations specific to each technology

You will not find exercises, quizzes, or simplified diagrams that pretend IPv6 is trivial. You will find technical depth.

How to Use This Book

Read it in order if you want a structured learning journey. The early chapters on addressing and Neighbor Discovery lay the foundation for everything that follows. You cannot securely firewall IPv6 if you do not understand that ICMPv6 is not optional. You cannot troubleshoot routing if you do not understand how OSPFv3 differs from OSPFv2.

Use it as a reference if you already have IPv6 deployed and need specific information. Each chapter stands alone enough to be useful when you are designing a dual-stack campus network, configuring NAT64 for IPv6-only data center access, or writing Ansible playbooks to manage IPv6 across hundreds of servers.

The configuration examples are organized by platform. Each example block identifies the platform, relevant version or context assumptions, prerequisite conditions, the exact configuration or commands, and verification steps. Adapt them to your environment. Test them in a lab first. Production networks are unforgiving of IPv6 misconfiguration.

Several chapters contain decision frameworks that help you choose between alternatives. Should you use SLAAC or DHCPv6? Stateless or stateful NAT64? 464XLAT or dual-stack for mobile access? These decisions are not

purely technical. They involve operational tradeoffs, vendor support considerations, and organizational capabilities. The book presents the factors so you can make informed choices for your environment.

One more note about notation. IPv6 addresses in this book follow RFC 5952 standard conventions: lowercase hex digits, longest run of zero groups compressed with ::, no leading zeros in groups [4]. When addresses are used in documentation examples, the 2001:db8::/32 prefix is used as specified in RFC 3849 [5].

The rest of this book will help you design, deploy, secure, operate, and troubleshoot IPv6 networks at production scale. Let us begin.

Chapter 1: IPv6 Addressing Fundamentals and Design Principles

IPv6 Address Format and Notation

An IPv6 address is 128 bits, four times the length of IPv4. This is not simply more of the same. The expanded address space enables new architectural possibilities: hierarchical aggregation that shrinks global routing tables, per-device globally routable addresses without NAT, and address ranges dedicated to specific purposes that have no IPv4 equivalent.

RFC 4291 defines the addressing architecture [6]. Addresses are written in textual form as eight groups of four hexadecimal digits separated by colons. For example:

```
2001:0db8:0000:0000:0000:ff00:0042:8329
```

RFC 5952 specifies two abbreviation rules that must be followed for interoperable representation [7]:

1. Leading zeros within each group may be omitted: `2001:db8:0:0:0:ff00:42:8329`
2. One consecutive sequence of zero groups may be replaced with `::`, but only once per address: `2001:db8::ff00:42:8329`

When both rules apply, the result is the standard compressed form. The `::` compression must represent the longest sequence of consecutive all-zero groups. When two sequences are equal length, compress the leftmost. Never produce ambiguous addresses.

Prefixes are written with a slash and prefix length: `2001:db8:1000::/36`. This notation means the first 36 bits are the network prefix and the remaining 92 bits identify interfaces within that prefix.

The full address space is divided into blocks by the most significant bits. RFC 3513 originally defined the Global Unicast Address space as `2000::/3`, covering

addresses from `2000::` to `3fff:ffff:ffff:ffff:ffff:ffff:ffff:ffff`. That space has been further allocated to Regional Internet Registries. For practical purposes, any address starting with 2 or 3 in the first hex digit is a global unicast address.

Several special-purpose blocks serve specific roles. The link-local range `fe80::/10` is for addresses valid only on a single physical or logical link. Unique Local Addresses `fc00::/7` are the IPv6 equivalent of private IPv4 addresses but with collision avoidance built in. The documentation prefix `2001:db8::/32` is reserved for examples and must never appear in production [8].

Understanding the structure of these ranges is critical because they determine routing behavior, firewall policy, and host configuration. An address that looks random is not. Its most significant bits tell you exactly what type of address it is and how it should be treated.

Address Types: Unicast, Multicast, Anycast

IPv6 defines three address types, each with distinct semantics.

Unicast addresses identify a single interface. A packet sent to a unicast address is delivered to exactly one interface. Unicast is by far the most common address type. IPv6 unicast addresses are subdivided into global unicast, link-local, unique local, and several special-purpose categories. RFC 4291 specifies the structure and assignment rules [9].

Multicast addresses identify a group of interfaces. A packet sent to a multicast address is delivered to all interfaces that have joined that multicast group. Unlike IPv4 broadcast, which floods an entire link whether receivers are present or not, IPv6 has no broadcast at all. Multicast replaces broadcast entirely. This matters for protocols. IPv6 Neighbor Discovery uses multicast addresses like `ff02::1` for all nodes on a link and `ff02::2` for all routers on a link. Every IPv6 node must support multicast.

Multicast addresses have the format `ff00::/8`. The next four bits specify scope: 1 for interface-local, 2 for link-local, 5 for site-local, 8 for organization-local, and e for global. The scope field is critical because it determines how far multicast traffic may propagate. A link-local multicast address never leaves the local link.

Anycast addresses are assigned to multiple interfaces but a packet is delivered to only one of them: the nearest one according to routing metrics.

Anycast is not a special address format. It is a routing configuration choice. You assign the same address to multiple routers in different locations. Routing protocols advertise that address from each location. Packets follow the shortest path to whichever instance is closest.

Anycast is powerful for distributed services. DNS root servers use anycast. Content delivery networks use anycast. NAT64 gateways can be deployed with anycast for redundancy and load distribution. The key operational requirement is stateless service: the server must not maintain per-client state because subsequent packets from the same client may arrive at different anycast instances.

Link-Local and Unique Local Addresses

Link-local addresses are fundamental to IPv6 operation. Every IPv6-enabled interface must have a link-local address. The prefix is `fe80::/10`. The next 54 bits are zero. The final 64 bits are the interface identifier, which can be derived from the MAC address using EUI-64, generated randomly per RFC 4941, or generated stably per RFC 7217 [10,11].

Link-local addresses serve several essential functions:

- Neighbor Discovery uses link-local addresses as source addresses
- Router Solicitation messages are sent from the unspecified address `::` but replies use link-local
- Default gateways are identified by their link-local addresses, not global addresses
- Point-to-point links can operate with link-local addresses only

Never filter all link-local traffic. Doing so breaks IPv6 fundamentally. Link-local addresses are never routed beyond the local link by standard routers. They are automatically scoped.

Unique Local Addresses are defined in RFC 4193 [12]. They occupy the `fc00::/7` range. The format is:

```
fc00:/7 | 1 | 40-bit Global ID | 16-bit Subnet ID | 64-bit Interface ID
```

The first bit after the prefix is set to 1 for randomly generated Global IDs: `fd00::/8`. The 40-bit Global ID is generated randomly to minimize collision risk across independent networks. RFC 4193 recommends using a cryptographically random number generator for this value [13].

ULAs are useful for internal addressing that remains stable regardless of external address changes. They work well in data center fabrics, for interconnect links between devices that do not need external reachability, and as a stable addressing layer that survives ISP changes. They are not a replacement for global addresses for Internet-facing services. They are a complement.

The key tradeoff is operability versus reachability. ULAs simplify internal operations because addresses never change. They also create complexity when you need to reach ULA-addressed hosts from outside your network. There is no global routing for ULA space. Tunneling or translation is required. Plan accordingly.

Global Unicast Addressing and Prefix Hierarchies

Global Unicast Addresses form the backbone of IPv6 Internet connectivity. They are allocated from the `2000::/3` space by IANA to Regional Internet Registries, then by RIRs to Local Internet Registries and large end sites, then by LIRs to customers.

The global routing table grows linearly with the number of non-aggregatable prefixes. That is a fundamental constraint. To keep routing table size manageable, the IPv6 hierarchy is designed for aggregation at every level. An RIR allocates large blocks such as `/23` or `/22` to an LIR. The LIR allocates `/32` or `/36` prefixes to enterprises or large customers. Those enterprises subdivide internally with perfectly aggregatable sub-prefixes.

For an enterprise receiving a `/36`, the internal hierarchy might look like this:

- Organization prefix: `2001:db8:1000::/36`
- Site A: `2001:db8:1000::/48`
  - Data center floor 1: `2001:db8:1000:1000::/56`
    - Server subnet 1: `2001:db8:1000:1000::/64`
    - Server subnet 2: `2001:db8:1000:1001::/64`
  - Corporate LAN: `2001:db8:1000:2000::/56`
    - VLAN 10: `2001:db8:1000:2000::/64`

        * VLAN 20: `2001:db8:1000:2001::/64`
- Site B: `2001:db8:2000::/48`

Each level allocates a contiguous block that can be aggregated by routing protocols. A site edge router advertises a `/48` to the provider. The provider aggregates all customer `/48`s into the `/36` or `/32` they received. The result is efficient routing table utilization.

The /64 subnet is the standard size for IPv6 LAN segments. It is not a suggestion. RFC 4291 and RFC 4862 require /64 subnets for SLAAC to function [14]. Neighbor Discovery expects /64 boundaries. Routing protocols optimize for /64. Use /64 for broadcast-domain-equivalent segments unless you have a specific reason and understand the implications.

Shorter prefixes than /64 are allowed for point-to-point links. RFC 6177 permits prefixes as short as /127 for point-to-point links between two nodes [15]. Using a /127 eliminates the need for proxy NDP on point-to-point links and reduces address space consumption. Not all platforms support /127. Check your hardware and software before deploying.

Longer prefixes than /64 for multi-access links are strongly discouraged. They break SLAAC, complicate Neighbor Discovery, and create compatibility problems with many implementations. If you absolutely must use a non-/64 subnet on a LAN, ensure that all connected devices support the prefix length and that you are not relying on SLAAC.

Addressing Design Principles and Pitfalls

Good IPv6 addressing design follows several principles.

First, design for aggregation. Every prefix you assign should be aggregatable with its siblings. Avoid scattered allocations that force route leaks. If you have a `/48` for Site A, subdivide it into `/56`s and `/64`s that are contiguous. Number them sequentially. This allows your site router to advertise a single aggregated prefix instead of dozens of more-specific routes.

Second, be consistent. Use the same numbering scheme across all sites. If Site A uses `::/56` for server subnets and Site B uses `::/56` for server subnets, operators can recognize subnet purposes by inspection. This consistency reduces configuration errors and speeds troubleshooting.

Third, document everything. IPv6 addresses are harder for humans to read than IPv4 addresses. A well-maintained IP address management system is not optional for serious IPv6 deployments. NetBox, IPAM from Infoblox, or even a carefully maintained spreadsheet is better than nothing. Record which prefix belongs to which subnet, which VLAN, which physical switch, and which building.

Fourth, separate management and data traffic. Use distinct subnets for device management and for data plane traffic. This allows granular firewall policies. Management traffic can be restricted to a specific IPv6 range. Data traffic flows through other ranges.

Fifth, plan for growth even though addresses are abundant. You have billions of addresses per /48. You will never run out. But you have a limited number of subnets if you allocate /64s from a /48. A /48 gives you 65536 /64 subnets. A /56 gives you 256. Plan your subnet hierarchy accordingly. A common pattern is to allocate /48 per site, /56 per VLAN or physical segment group, and /64 per actual subnet.

Common pitfalls to avoid:

- Using the deprecated site-local prefix `fec0::/10`. RFC 3879 deprecated it and it causes routing problems [16]. Use unique local addresses `fc00::/7` instead.
- Subnetting a /64 into smaller pieces on a LAN. This breaks SLAAC and confuses many implementations.
- Filtering ICMPv6 types required for Neighbor Discovery. IPv6 will silently fail if you block Router Advertisements or Neighbor Solicitations.
- Relying solely on privacy extension addresses for servers. Privacy addresses change periodically. Servers need stable addresses for DNS records and certificates.
- Forgetting that link-local addresses are mandatory. Every IPv6 interface needs one. Design around that requirement.
- Using 2001:db8:: addresses in production. That prefix is for documentation only and is not routable on the public Internet.

The addressing foundation you build now will affect your IPv6 deployment for years. Get the hierarchy right. Document it. Enforce consistency. Everything else becomes easier.

Chapter 2: IPv6 Address Planning and Allocation Strategy

Understanding RIR Allocation and Delegation

Before you plan your internal IPv6 addressing, you need to understand how address space reaches you from the top down. The IANA allocates large blocks from the global unicast space `2000::/3` to five Regional Internet Registries: ARIN for North America, RIPE NCC for Europe and the Middle East, APNIC for Asia-Pacific, LACNIC for Latin America, and AFRINIC for Africa.

Each RIR allocates to Local Internet Registries such as ISPs and large enterprises. End-user organizations receive a prefix from their ISP or directly from their RIR if they qualify as a Large End User Site. The minimum allocation varies by RIR but a typical enterprise allocation is `/36` from the RIR or `/48` from an ISP.

The allocation structure follows routing aggregation principles. An ISP might receive `/24` from the RIR and delegate `/36` prefixes to enterprise customers and `/56` prefixes to residential customers. This ensures that each customer prefix aggregates into the ISP's larger block, which in turn aggregates into the RIR's block.

RFC 7046 recommends that ISPs assign at least `/48` to any customer that needs to subnet internally [17]. For residential customers, RFC 7084 recommends `/56` as the standard allocation size, providing 256 /64 subnets per home or small office [18]. A /56 is the current sweet spot for residential broadband because most home routers can only create a limited number of subnets.

When you request an allocation from an ISP, specify your needs clearly. A single-site organization with one LAN needs a `/48`. A multi-site organization should request a `/40` or larger from the RIR directly, or negotiate a larger prefix with an ISP. The key question is whether you need to aggregate routes across multiple sites. If yes, you need enough prefix space to assign distinct, aggregatable prefixes to each site.

IPv6 prefix delegation is the mechanism that makes this work dynamically. DHCPv6 Prefix Delegation allows an ISP to assign a prefix to a customer's router, which then sub-delegates sub-prefixes to internal networks. This is essential for multi-homed sites that receive prefixes from different ISPs and need to advertise their own prefixes into BGP.

Designing Your Internal Address Hierarchy

A well-designed internal IPv6 hierarchy makes routing efficient, firewall policies manageable, and troubleshooting straightforward. The fundamental goal is to ensure that every prefix you assign is aggregatable with its siblings.

Start from the top. If you have a `/36`, allocate contiguous `/40`s or `/48`s to major organizational units. A common pattern is one `/48` per physical site:

```
Organization prefix: 2001:db8:1000::/36

Site: New York           2001:db8:1000::/48
Site: London             2001:db8:2000::/48
Site: Tokyo              2001:db8:3000::/48
Site: Frankfurt          2001:db8:4000::/48
```

Within each site `/48`, allocate `/56`s for functional groups:

```
Site: New York           2001:db8:1000::/48
  Data Center            2001:db8:1000:1000::/56
  Corporate LAN          2001:db8:1000:2000::/56
  Guest WiFi             2001:db8:1000:3000::/56
  Management             2001:db8:1000:4000::/56
  IoT devices            2001:db8:1000:5000::/56
```

Within each `/56`, allocate `/64`s for individual VLANs or physical segments:

```
Data Center:             2001:db8:1000:1000::/56
  Server VLAN 10         2001:db8:1000:1000::/64
  Server VLAN 20         2001:db8:1000:1001::/64
  Database VLAN 30       2001:db8:1000:1002::/64
  Storage VLAN 40        2001:db8:1000:1003::/64
  Management VLAN 50     2001:db8:1000:1004::/64
```

This hierarchy has several advantages. The site edge router advertises a single `/48` to upstream providers. Internal core routers aggregate `/56`s. Access switches route between `/64`s. Each level aggregates its children into a single prefix.

Numbering within a hierarchy should be consistent and intentional. Many organizations use the second or third hextet for site identification and subsequent hextets for functional group identification. For example, `2001:db8:SITE:GROUP:LAN::/64` makes the purpose of each subnet immediately recognizable.

RIPE has published guidance in RIPE-690 on IPv6 prefix assignment for end-users [19]. The document recommends a `/48` for any end-user that needs to create multiple subnets and strongly discourages delegating prefixes longer than `/56`.

Prefix Sizing: /48, /56, /64, and When to Use Each

The right prefix size depends on the context. There are standard conventions for good reasons.

/48 is the standard allocation size for enterprises and large sites. It provides 65,536 /64 subnets, which is more than enough for most organizations. If you receive a /48 from your ISP or RIR, subdivide it into /56s for departments or physical segments and /64s for individual VLANs or LAN segments.

/56 is the standard allocation for residential customers and small offices. It provides 256 /64 subnets, sufficient for a home network with separate subnets for wired devices, guest WiFi, IoT devices, and smart home systems. RFC 7084 specifies /56 as the recommended size for residential networks [20].

/64 is mandatory for LAN segments that use SLAAC. RFC 4862 requires /64 prefixes for stateless address autoconfiguration to work correctly [21]. Neighbor Discovery Protocol assumes /64 boundaries for solicited-node multicast address calculation. Routing protocols and firewalls optimize for /64 subnets.

Use /64 for every broadcast-domain-equivalent segment unless you have a specific reason to do otherwise.

Smaller than /64 is acceptable for point-to-point links between exactly two devices. RFC 6177 permits prefixes as small as /127 for point-to-point links [22]. Using a /127 on a point-to-point link conserves address space and eliminates proxy NDP requirements. Some platforms only support /126 or /128 for point-to-point links. Check your equipment.

Larger than /64 for LAN segments is problematic. It breaks SLAAC, complicates Neighbor Discovery, and may not work with all implementations. Avoid it unless you have no other option and understand the consequences.

A practical sizing guideline for enterprise networks:

- Organization: /36 or /40 from RIR
- Each site: /48
- Each major functional group at a site: /56
- Each VLAN or physical segment: /64
- Each point-to-point link: /127 or /126

For data center fabrics, consider allocating a /48 or /44 to the data center and subdividing by pod or rack group with /56s. This gives you flexibility to add segments without renumbering.

Special-Purpose and Reserved Prefixes

IPv6 defines numerous special-purpose prefixes for specific functions. Knowing them is essential for correct configuration and security policy. The IANA maintains the authoritative registry at https://www.iana.org/assignments/iana-ipv6-special-registry [23].

Key prefixes you will encounter:

- `::/128` - Unspecified address. Used as source address when sending a Router Solicitation. Never appear in routing tables or firewall rules as a valid source.
- `::1/128` - Loopback address. Equivalent to 127.0.0.1.
- `::/96` - IPv4-mapped IPv6 addresses. Legacy format used in some dual-stack applications.

- `::ffff:0:0/96` - IPv4-compatible IPv6 addresses. Deprecated.
- `100::/64` - Discard prefix. Packets to this range are silently dropped.
- `2001::/23` - Teredo tunneling. Largely obsolete.
- `2001:db8::/32` - Documentation prefix. Never use in production.
- `2002::/16` - 6to4. Deprecated for new deployments.
- `fc00::/7` - Unique Local Addresses. fc00::/8 is reserved, fd00::/8 is used for randomly generated IDs.
- `fe80::/10` - Link-local addresses. Mandatory on every IPv6 interface.
- `ff00::/8` - Multicast addresses.

RFC 5156 compiles special-use IPv6 addresses for use as a checklist for route filtering [24]. Implement your ingress filtering policies to drop routes and packets with these special prefixes unless you explicitly need them.

For ingress filtering at your network edge, implement RPKI where possible and supplement with bogon filtering that includes these special prefixes. At minimum, filter:

- `::/128` as a source (unless it is a Router Solicitation on a local link)
- `::/96` as a source from the Internet (these are internal)
- `fc00::/7` from the public Internet
- `fe80::/10` from the public Internet
- `ff00::/8` as a source address

The documentation prefix `2001:db8::/32` should also be filtered at your perimeter. Legitimate Internet traffic should never use this prefix.

Documenting and Automating Address Assignment

IPv6 addresses are long and opaque. A human cannot memorize the structure of your address plan from memory. Documentation is not optional. Without it, you will make mistakes, create conflicts, and waste time troubleshooting problems that would be instantly visible with a clear plan.

Use an IP Address Management system that supports IPv6. NetBox is a popular open-source choice that provides full IPv6 support, hierarchy visualization, API access for automation, and integration with configuration management tools [25]. Commercial options include Infoblox IPAM and Cisco IP Services Manager.

Your IPAM system should track:

- Aggregates and prefixes at every level
- Which device, VLAN, or network segment owns each prefix
- Individual IP addresses assigned to servers or critical infrastructure
- DNS names associated with addresses
- VRF or routing instance assignments
- Expiration dates for temporary allocations

Automate address assignment from your IPAM system. Use the NetBox API with Ansible to pull address assignments and push configurations to network devices. Use Terraform to allocate addresses in cloud environments and record them in your IPAM system.

Here is a pattern for automating with NetBox and Ansible:

Platform: NetBox with Ansible, using netbox.netbox collection. Prerequisites: NetBox installed, API token created with read/write access, Ansible control node with collection installed.

```yaml
# netbox-ipv6-prefixes.yml
---
- name: Create IPv6 prefix hierarchy in NetBox
  hosts: localhost
  gather_facts: false
  collections:
    - netbox.netbox
  vars:
    netbox_url: "https://netbox.example.com"
    netbox_token: "{{ lookup('env', 'NETBOX_TOKEN') }}"
    org_prefix: "2001:db8:1000::/36"

  tasks:
    - name: Create organization aggregate
      netbox_prefix:
        netbox_url: "{{ netbox_url }}"
        netbox_token: "{{ netbox_token }}"
        data:
          prefix: "{{ org_prefix }}"
          site: null
          description: "Organization IPv6 allocation"
          status: "reserved"
        state: present

    - name: Create site prefixes
      netbox_prefix:
        netbox_url: "{{ netbox_url }}"
        netbox_token: "{{ netbox_token }}"
```

```
29          data:
30            prefix: "{{ item.prefix }}"
31            site: "{{ item.site }}"
32            description: "{{ item.description }}"
33            status: "active"
34            parent: "{{ org_prefix }}"
35          state: present
36        loop:
37          - { site: "new-york", prefix: "2001:db8:1000::/48", description: "New
            ↪  York site" }
38          - { site: "london", prefix: "2001:db8:2000::/48", description: "London
            ↪  site" }
39          - { site: "tokyo", prefix: "2001:db8:3000::/48", description: "Tokyo
            ↪  site" }
```

This approach keeps your IPAM system as the source of truth. Network device configurations are generated from it, ensuring consistency across the infrastructure. When you need to add a subnet, you update NetBox and run the playbook. The result is repeatable, auditable, and less error-prone than manual address management.

For cloud environments, use provider tools to allocate IPv6 CIDRs and integrate them into your IPAM system. AWS provides /56 blocks for VPCs. Azure allows custom /48 or larger address spaces. GCP provides internal /48 ULAs and external IPv6 ranges. Record all allocations in your IPAM system regardless of the source.

The effort spent building a proper documentation and automation framework pays dividends every time you need to troubleshoot a routing problem, expand a subnet, or audit your address space. Do it from the start.

Chapter 3: Stateless Address Autoconfiguration (SLAAC)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

The SLAAC Process: Step by Step

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

EUI-64 and Privacy Extensions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Router Advertisement Parameters

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

SLAAC Security Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

When to Use SLAAC and When Not To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 4: DHCPv6: Stateful and Stateless Address Assignment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DHCPv6 Architecture and Message Flows

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Stateful vs. Stateless DHCPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DHCPv6 Options and Extended Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DHCPv6 Security: Guards and Authentication

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DHCPv6 Design Patterns for Enterprises

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 5: Router Advertisements and Prefix Delegation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Router Advertisement Structure and Fields

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Tuning RA Parameters for Your Network

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Prefix Delegation (DHCPv6-PD)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

RA Guard and Security Extensions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

RA Design for Multi-Homed Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 6: Neighbor Discovery Protocol (NDP)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

NDP Message Types and Operation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Neighbor Solicitation and Advertisement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Duplicate Address Detection

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Neighbor Cache Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

NDP Security and Spoofing Defenses

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 7: ICMPv6 and Path MTU Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

ICMPv6 Type and Code Structure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Essential ICMPv6 Messages

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Path MTU Discovery and Black Holes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Fragmentation in IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Filtering ICMPv6 Without Breaking the Network

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 8: DNS and IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

AAAA Records and IPv6 Reverse Lookup

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNS Zone Design for Dual-Stack Environments

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNSSEC and IPv6 Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNS64 and Synthetic AAAA Records

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNS Troubleshooting for IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 9: IPv6 Routing Protocols and Route Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Static Routing for IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

OSPFv3: Protocol Mechanics and Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

RIPng: Operation and Limitations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

BGP and IPv6: MP-BGP Implementation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IS-IS for IPv6 and Large Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Route Design and Aggregation Strategy

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 10: Dual-Stack Architecture and Implementation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack Design Principles

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack on Linux

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack on Enterprise Network Equipment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack Application Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack Pitfalls and Troubleshooting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 11: IPv6-Only Network Architectures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

When IPv6-Only Makes Sense

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Designing an IPv6-Only Data Center

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Client Considerations and Compatibility

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Service Dependencies and Translation Needs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Real-World IPv6-Only Deployments

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 12: Transition Mechanisms: Tunneling and Brokering

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

6to4 and Automatic Tunneling

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

6rd and ISP-Scale Tunneling

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

ISATAP for Intra-Site Connectivity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Teredo and NAT Traversal

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Manual GRE and IPv6-in-IPv4 Tunnels

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

VXLAN and Overlays with IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Transition Mechanism Selection Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 13: Translation Mechanisms: NAT64, DNS64, and SIIT

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

NAT64 Architecture and Prefix Selection

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNS64 Operation and Integration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Stateless vs. Stateful NAT64

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

464XLAT and Carrier-Grade Deployment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Application-Layer Translation Challenges

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

NAT64/DNS64 Operational Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 14: IPv4/IPv6 Interoperability Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Happy Eyeballs and DNS Resolution Ordering

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Dual-Stack Service Deployment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Application Compatibility Checklist

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

API and Service Mesh Implications

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Interoperability Testing Strategies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 15: Enterprise Networking with IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Campus and LAN IPv6 Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Enterprise WAN and Branch IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 for Wireless Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 for VoIP and Unified Communications

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 for Enterprise IoT and OT

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 16: Data Center and Cloud IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Data Center Fabric IPv6 Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 in VMware and Hyper-V

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 for Containers and Docker

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Kubernetes and IPv6 Networking

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Public Cloud IPv6: AWS, Azure, GCP

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 17: Operating Systems and Host IPv6 Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Linux IPv6 Configuration and Tuning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Windows Server and Client IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

macOS and BSD IPv6 Notes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Web Servers: Apache, Nginx with IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Mail and Database Servers IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 18: Firewalls and IPv6 Security Architecture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Security Model and Differences from IPv4

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Firewall Design for IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 ACL Design and Implementation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Zone-Based Firewalls and IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Next-Gen Firewalls and IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 19: IPv6 Security: IPsec, VPNs, and Segmentation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPsec Integration with IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Site-to-Site VPNs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Remote Access and SSL VPNs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Microsegmentation in IPv6 Environments

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6-Specific Attack Surfaces

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 20: Multicast, Anycast, and Service Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Multicast Addressing and Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

PIM and Multicast Routing with IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Anycast: Design and Deployment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Multicast DNS and Service Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IoT and mDNS in IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 21: IPv6 Troubleshooting and Packet Analysis

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Troubleshooting Methodology

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Packet Capture with tcpdump and Wireshark

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Analyzing NDP and RA Problems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Routing Troubleshooting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Common Failure Modes and Fixes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 22: IPv6 Observability, Monitoring, and Performance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Monitoring: What to Measure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

NetFlow, sFlow, and IPv6 Flow Data

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Logging and SIEM for IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Performance Optimization and MTU Tuning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Alerting and Capacity Planning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 23: Automation and Infrastructure as Code for IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Ansible for IPv6 Configuration Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Terraform and IPv6 in Cloud/Network

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Python and Netmiko for IPv6 Automation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPAM with NetBox and IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

CI/CD for IPv6 Network Changes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Chapter 24: IPv6 Migration Strategies and Real-World Deployment Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Planning the IPv6 Migration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Migration Approach: Dual-Stack First

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Migration Approach: Site-by-Site Deployment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Migration Approach: Network Segment by Segment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Migration Approach: IPv6-Only with NAT64/DNS64

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Real-World Deployment Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Common Migration Mistakes and How to Avoid Them

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Migration Timeline and Expectations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Conclusion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Core IPv6 Standards

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

ICMPv6 and PMTUD

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Addressing and Allocation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DHCPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Privacy Extensions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

DNS and IPv6

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Transition and Migration Mechanisms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Routing Protocols

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Operational and Deployment Guidance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

IPv6 Adoption Statistics

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Tools and Platforms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.

Additional References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/advancedipv6engineering.