



Troubleshooting BGP

A Practical Guide to Understanding
and Troubleshooting BGP

Exclusive Offer – 40% OFF

Cisco Press Video Training

livelessons®

ciscopress.com/video

Use coupon code **CPVIDEO40** during checkout.



Video Instruction from Technology Experts



Advance Your Skills

Get started with fundamentals, become an expert, or get certified.



Train Anywhere

Train anywhere, at your own pace, on any device.



Learn

Learn from trusted author trainers published by Cisco Press.

Try Our Popular Video Training for FREE!

ciscopress.com/video

Explore hundreds of **FREE** video lessons from our growing library of Complete Video Courses, LiveLessons, networking talks, and workshops.

Cisco Press

ciscopress.com/video

Troubleshooting BGP: A Practical Guide to Understanding and Troubleshooting BGP

Table of Contents

Cover

Title Page

Copyright Page

About the Authors

Acknowledgments

Contents

Foreword

Introduction

Part I: BGP Fundamentals

Chapter 1 BGP Fundamentals

Border Gateway Protocol

Autonomous System Numbers

Path Attributes

Loop Prevention

Address Families

BGP Sessions

Inter-Router Communication

BGP Messages

OPEN

Hold Time

BGP Identifier

KEEPALIVE

UPDATE

Table of Contents

NOTIFICATION Message

BGP Neighbor States

Idle

Connect

Active

OpenSent

OpenConfirm

Established

Basic BGP Configuration

IOS

IOS XR

NX-OS

Verification of BGP Sessions

Prefix Advertisement

BGP Best-Path Calculation

Route Filtering and Manipulation

IBGP

IBGP Full Mesh Requirement

Peering via Loopback Addresses

EBGP

EBGP and IBGP Topologies

Next-Hop Manipulation

IBGP Scalability

Route Reflectors

Loop Prevention in Route Reflectors

Out-of-Band Route Reflectors

Confederations

BGP Communities

Route Summarization

Aggregate-Address

Flexible Route Suppression

Selective Prefix Suppression

Table of Contents

- Leaking Suppressed Routes
- Atomic Aggregate
- Route Aggregation with AS_SET
- Route Aggregation with Selective Advertisement of AS-SET
- Default Route Advertisement
- Default Route Advertisement per Neighbor

Remove Private AS

Allow AS

LocalAS

Summary

References

Part II: Common BGP Troubleshooting

Chapter 2 Generic Troubleshooting Methodologies

Identifying the Problem

Understanding Variables

Reproducing the Problem

- Setting Up the Lab

- Configuring Lab Devices

- Triggering Events

Sniffer-Packet Capture

- SPAN on Cisco IOS

- SPAN on Cisco IOS XR

- SPAN on Cisco NX-OS

- Remote SPAN

Platform-Specific Packet Capture Tools

- Netdr Capture

- Embedded Packet Capture

- Ethalyzer

Logging

Event Monitoring/Tracing

Summary

Table of Contents

Reference

Chapter 3 Troubleshooting Peering Issues

BGP Peering Down Issues

Verifying Configuration

Verifying Reachability

Find the Location and Direction of Packet Loss

Verify Whether Packets Are Being Transmitted

Use Access Control Lists to Verify Whether Packets Are Received

Check ACLs and Firewalls in Path

Verify TCP Sessions

Simulate a BGP Session

Demystifying BGP Notifications

Decode BGP Messages

Troubleshoot Blocked Process in IOS XR

Verify BGP and BPM Process State

Verify Blocked Processes

Restarting a Process

BGP Traces in IOS XR

BGP Traces in NX-OS

Debugs for BGP

Troubleshooting IPv6 Peers

Case StudySingle Session Versus Multisession

Multisession Capability

Single-Session Capability

BGP Peer Flapping Issues

Bad BGP Update

Hold Timer Expired

Interface Issues

Physical Connectivity

Physical Interface

Input Hold Queue

TCP Receive Queue

MTU Mismatch Issues

Table of Contents

High CPU Causing Control-Plane Flaps

Control Plane Policing

CoPP on NX-OS

Local Packet Transport Services

Dynamic BGP Peering

Dynamic BGP Peer Configuration

Dynamic BGP Challenges

Misconfigured MD5 Password

Resource Issues in a Scaled Environment

TCP Starvation

Summary

References

Chapter 4 Troubleshooting Route Advertisement and BGP Policies

Troubleshooting BGP Route Advertisement

Local Route Advertisement Issues

Route Aggregation Issues

Route Redistribution Issues

BGP Tables

Receiving and Viewing Routes

Troubleshooting Missing BGP Routes

Next-Hop Check Failures

Bad Network Design

Validity Check Failure

AS-Path

Originator-ID/Cluster-ID

BGP Communities

BGP Communities: No-Advertise

BGP Communities: No-Export

BGP Communities: Local-AS (No Export SubConfed)

Mandatory EBGP Route Policy for IOS XR

Filtering of Prefixes by Route Policy

Conditional Matching

Table of Contents

Access Control Lists (ACL)

Prefix Matching

Regular Expressions (Regex)

UnderScore _

Caret ^

Dollar Sign \$

Brackets []

Hyphen -

Caret in Brackets [^]

Parentheses () and Pipe |

Period

Plus Sign +

Question Mark ?

Asterisk *

Looking Glass and Route Servers

Conditionally Matching BGP Communities

Troubleshooting BGP Router Policies

IOS and NX-OS Prefix-Lists

IOS and NX-OS AS-Path ACLs

Route-Map Processing

IOS and NX-OS Route-Maps

IOS XR Route-Policy Language

Incomplete Configuration of Routing Policies

Conditional BGP Debugs

Summary

Further Reading

References in This Chapter

Chapter 5 Troubleshooting BGP Convergence

Understanding BGP Route Convergence

BGP Update Groups

BGP Update Generation

Troubleshooting Convergence Issues

Table of Contents

- Faster Detection of Failures
- Jumbo MTU for Faster Convergence
- Slow Convergence due to Periodic BGP Scan
- Slow Convergence due to Default Route in RIB
- BGP Next-Hop Tracking
- Selective Next-Hop Tracking
- Slow Convergence due to Advertisement Interval
- Computing and Installing New Path
- Troubleshooting BGP Convergence on IOS XR
- Verifying Convergence During Initial Bring Up
- Verifying BGP Reconvergence in Steady State Network
- Troubleshooting BGP Convergence on NX-OS

BGP Slow Peer

- BGP Slow Peer Symptoms
- High CPU due to BGP Router Process
- Traffic Black Hole and Missing Prefixes in BGP table
- BGP Slow Peer Detection
- Verifying OutQ value
- Verifying SndWnd
- Verifying Cache Size and Pending Replication Messages
- Workaround
- Changing Outbound Policy
- Advertisement Interval
- BGP Slow Peer Feature
- Static Slow Peer
- Dynamic Slow Peer Detection
- Slow Peer Protection
- Slow Peer Show Commands

Troubleshooting BGP Route Flapping

Summary

Reference

Part III: BGP Scalability Issues

Table of Contents

Chapter 6 Troubleshooting Platform Issues Due to BGP

Troubleshooting High CPU Utilization due to BGP

Troubleshooting High CPU due to BGP on Cisco IOS

High CPU due to BGP Scanner Process

High CPU due to BGP Router Process

High CPU Utilization due to BGP I/O Process

Troubleshooting High CPU due to BGP on IOS XR

Troubleshooting High CPU due to BGP on NX-OS

Capturing CPU History

Troubleshooting Sporadic High CPU Condition

Troubleshooting Memory Issues due to BGP

TCAM Memory

Troubleshooting Memory Issues on Cisco IOS Software

Troubleshooting Memory Issues on IOS XR

Troubleshooting Memory Issues on NX-OS

Restarting Process

Summary

References

Chapter 7 Scaling BGP

The Impact of Growing Internet Routing Tables

Scaling Internet Table on Various Cisco Platforms

Scaling BGP Functions

Tuning BGP Memory

Prefixes

Managing the Internet Routing Table

Paths

Attributes

Tuning BGP CPU

IOS Peer-Groups

IOS XR BGP Templates

NX-OS BGP Peer Templates

BGP Peer Templates on Cisco IOS

Soft Reconfiguration Inbound Versus Route Refresh

Table of Contents

- Dynamic Refresh Update Group
- Enhanced Route Refresh Capability
- Outbound Route Filtering (ORF)
- Prefix-Based ORF
- Extended CommunityBased ORF
- BGP ORF Format
- BGP ORF Configuration Example
- Maximum Prefixes
- BGP Max AS
- BGP Maximum Neighbors

Scaling BGP with Route Reflectors

- BGP Route Reflector Clusters
- Hierarchical Route Reflectors
- Partitioned Route Reflectors
- BGP Selective Route Download
- Virtual Route Reflectors
- BGP Diverse Path
- Shadow Route Reflectors
- Shadow Sessions

Route Servers

Summary

References

Chapter 8 Troubleshooting BGP Edge Architectures

BGP Multihoming and Multipath

- Resiliency in Service Providers
- EBGP and IBGP Multipath Configuration
- EIBGP Multipath
- R1
- R2
- R3
- R4
- R5
- AS-Path Relax

Table of Contents

Understanding BGP Path Selection

- Routing Path Selection Longest Match
- BGP Best-Path Overview
- Weight
- Local Preference
- Locally Originated via Network or Aggregate Advertisement
- Accumulated Interior Gateway Protocol (AIGP)
- Shortest AS-Path
- Origin Type
- Multi-Exit Discriminator (MED)
- EBGP over IBGP
- Lowest IGP Metric
- Prefer the Oldest EBGP Path
- Router ID
- Minimum Cluster List Length
- Lowest Neighbor Address

Troubleshooting BGP Best Path

- Visualizing the Topology
- Phase IInitial BGP Edge Route Processing
- Phase IIBGP Edge Evaluation of Multiple Paths
- Phase IIIFinal BGP Processing State
- Path Selection for the Routing Table

Common Issues with BGP Multihoming

- Transit Routing
- Problems with Race Conditions
- Peering on Cross-Link
- Expected Behavior
- Unexpected Behavior
- Secondary Verification Methods of a Routing Loop
- Design Enhancements
- Full Mesh with IBGP
- Problems with Redistributing BGP into an IGP

Summary

Table of Contents

References

Part IV: Securing BGP

Chapter 9 Securing BGP

The Need for Securing BGP

Securing BGP Sessions

Explicitly Configured Peers

IPv6 BGP Peering Using Link-Local Address

BGP Session Authentication

BGP Pass Through

EBGP-Multihop

BGP TTL Security

Filtering

Protecting BGP Traffic Using IPsec

Securing Interdomain Routing

BGP Prefix Hijacking

S-BGP

IPsec

Public Key Infrastructure

Attestations

soBGP

Entity Certificate

Authorization Certificate

Policy Certificate

BGP SECURITY Message

BGP Origin AS Validation

Route Origination Authorization (ROA)

RPKI Prefix Validation Process

Configuring and Verifying RPKI

RPKI Best-Path Calculation

BGP Remote Triggered Black-Hole Filtering

BGP Flowspec

Configuring BGP Flowspec

Table of Contents

Summary

References

Part V: Multiprotocol BGP

Chapter 10 MPLS Layer 3 VPN (L3VPN)

MPLS VPNs

MPLS Layer 3 VPN (L3VPN) Overview

Virtual Routing and Forwarding

Route Distinguisher

Route Target

Multi-Protocol BGP (MP-BGP)

Network Advertisement Between PE and CE Routers

MPLS Layer 3 VPN Configuration

VRF Creation and Association

IOS VRF Creation

IOS XR VRF Creation

NX-OS VRF Creation

Verification of VRF Settings and Connectivity

Viewing VRF Settings and Interface IP Addresses

Viewing the VRF Routing Table

VRF Connectivity Testing Tools

MPLS Forwarding

BGP Configuration for VPNv4 and PE-CE Prefixes

IOS BGP Configuration for MPLS L3VPN

IOS XR BGP Configuration for MPLS L3VPN

NX-OS BGP Configuration for MPLS L3VPN

Verification of BGP Sessions and Routes

Troubleshooting MPLS L3VPN

Default Route Advertisement Between PE-CE Routers

Problems with AS-PATH

Suboptimal Routing with VPNv4 Route Reflectors

Troubleshooting Problems with Route Targets

MPLS L3VPN Services

Table of Contents

RT Constraints

MPLS VPN Label Exchange

MPLS Forwarding

Summary

References

Chapter 11 BGP for MPLS L2VPN Services

L2VPN Services

Terminologies

Virtual Private Wire Service

Interworking

Configuration and Verification

VPWS BGP Signaling

Configuration

Virtual Private LAN Service

Configuration

Verification

VPLS Autodiscovery Using BGP

VPLS BGP Signaling

Troubleshooting

Summary

References

Chapter 12 IPv6 BGP for Service Providers

IPv6 BGP Features and Concepts

IPv6 BGP Next-Hop

IPv6 Reachability over IPv4 Transport

IPv4 Routes over IPv6 Next-Hop

IPv6 BGP Policy Accounting

IPv6 Provider Edge Routers (6PE) over MPLS

6PE Configuration

6PE Verification and Troubleshooting

IPv6 VPN Provider Edge (6VPE)

IPv6-Aware VRF

Table of Contents

- 6VPE Next-Hop
- Route Target
- 6VPE Control Plane
- 6VPE Data Plane
- 6VPE Configuration
- 6VPE Control-Plane Verification
- 6VPE Data Plane Verification

Summary

References

Chapter 13 VxLAN BGP EVPN

Understanding VxLAN

- VxLAN Packet Structure
- VxLAN Gateway Types

VxLAN Overlay

- VxLAN Flood-and-Learn Mechanism
- Configuration and Verification
- Ingress Replication

Overview of VxLAN BGP EVPN

- Distributed Anycast Gateway
- ARP Suppression
- Integrated Route/Bridge (IRB) Modes
- Asymmetric IRB
- Symmetric IRB
- Multi-Protocol BGP
- Configuring and Verifying VxLAN BGP EVPN

Summary

References

Part VI: High Availability

Chapter 14 BGP High Availability

- BGP Graceful-Restart
- BGP Nonstop Routing

Table of Contents

Bidirectional Forwarding Detection

- Asynchronous Mode
- Asynchronous Mode with Echo Function
- Configuration and Verification
- Troubleshooting BFD Issues
- BFD Session Not Coming Up
- BFD Session Flapping

BGP Fast-External-Fallover

BGP Add-Path

BGP best-external

BGP FRR and Prefix-Independent Convergence

- BGP PIC Core
- BGP PIC Edge
- Scenario 1IP PE-CE Link/Node Protection on CE Side
- Scenario 2IP MPLS PE-CE Link/Node Protection for Primary/ Backup
- BGP Recursion Host

Summary

References

Part VII: BGP: Looking Forward

Chapter 15 Enhancements in BGP

Link-State Distribution Using BGP

- BGP-LS NLRI
- BGP-LS Path Attributes
- BGP-LS Configuration
- IGP Distribution
- BGP Link-State Session Initiation

BGP for Tunnel Setup

Provider Backbone Bridging: Ethernet VPN (PBB-EVPN)

- EVPN NLRI and Routes
- EVPN Extended Community
- EVPN Configuration and Verification

Summary

Table of Contents

References

Index