

HPE 5945 configuration errors and solutions

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: Dac 5945

Cat. No.: B1669747

[Get Quote](#)

HPE 5945 Technical Support Center

Welcome to the technical support center for the HPE 5945 switch series. This resource is designed for researchers, scientists, and drug development professionals who may encounter network configuration challenges during their experiments. Below you will find troubleshooting guides and Frequently Asked Questions (FAQs) to help you resolve common issues.

Troubleshooting Guides

This section provides detailed, question-and-answer based guides to address specific configuration errors you might encounter with the HPE 5945 switch.

Intelligent Resilient Framework (IRF) Configuration

Question: Why are my two HPE 5945 switches failing to form an IRF fabric?

Answer:

There are several potential reasons for IRF fabric formation failure. Here is a step-by-step troubleshooting guide:

Experimental Protocol: IRF Formation Troubleshooting

- **Verify Physical Connectivity:** Ensure the IRF ports on both switches are physically connected. Remember to use a crossover pattern: IRF port 1 on the first switch should be connected to IRF port 2 on the second switch.

- **Check Firmware Compatibility:** Inconsistent firmware versions between switches can prevent IRF formation. Verify that both switches are running the exact same firmware version. Upgrading the firmware to a version that fixes known IRF-related bugs may be necessary. For instance, some firmware versions have known issues with forming an IRF fabric using 100GbE QSFP28 ports.[\[1\]](#)
- **Review IRF Configuration:**
 - **Member Priority:** The switch you intend to be the master should have a higher IRF member priority.
 - **IRF Port Binding:** Ensure that the physical ports are correctly bound to the logical IRF ports.
 - **Domain ID:** Both switches must be configured with the same IRF domain ID.
- **Execute Diagnostic Commands:** Use the `display irf configuration` command to review the current IRF settings on both switches. The `display irf topology` command can help visualize the current state of the IRF fabric.
- **Check for Module-Specific Restrictions:** Certain line cards, like the JH405A module, have port grouping restrictions for IRF. The IRF ports must be configured on either odd or even-numbered ports, not consecutive ones.

Question: What is an "IRF split-brain" scenario and how do I recover from it?

Answer:

An IRF split-brain occurs when the links between IRF members fail, causing the single logical fabric to split into two or more independent fabrics that are still active and using the same IP address. This can lead to network instability and connectivity issues.[\[2\]](#)

Recovery Protocol: IRF Split-Brain

- **Detection:** A Multi-Active Detection (MAD) mechanism, such as LACP, BFD, or ARP, is crucial for detecting a split-brain scenario.[\[1\]](#)[\[2\]](#)[\[3\]](#) When a split is detected, the MAD protocol

will put the IRF fabric with the lower-priority master into a "Recovery" state, shutting down most of its interfaces to prevent IP address conflicts.

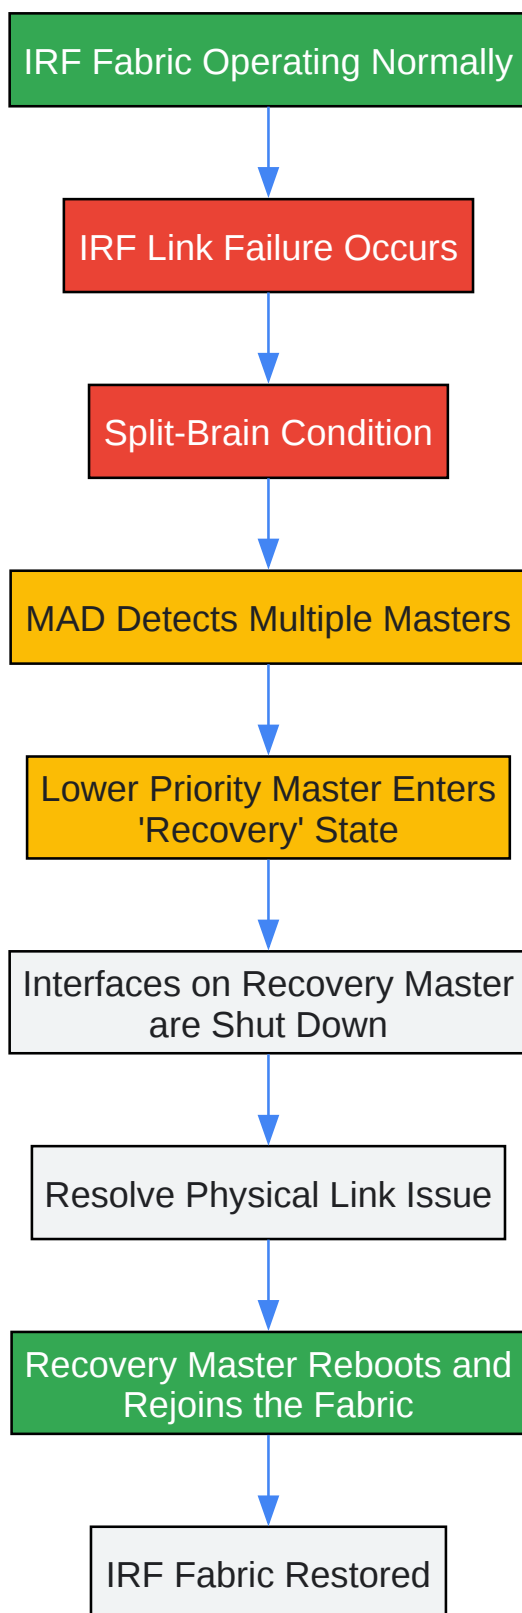
- **Identify the Cause:** Investigate the reason for the IRF link failure. This could be a faulty cable, a malfunctioning transceiver, or a misconfiguration.
- **Resolve the Link Failure:** Replace any faulty hardware or correct the configuration that caused the link to go down.
- **Merge the Fabrics:** Once the physical link is restored, the IRF fabric in the "Recovery" state will typically reboot and rejoin the active fabric automatically.
- **Manual Recovery:** If the active fabric fails before the link is restored, you can use the `mad restore` command on the inactive fabric to bring its interfaces back up and restore services.

Question: I'm observing MAC address flapping after an IRF split. Why is this happening?

Answer:

This can be expected behavior if you are using BFD for Multi-Active Detection (MAD) and have `irf mac-address persistent` always configured. In a split-brain scenario, both masters may attempt to use the same MAC address, leading to flapping on upstream devices. To mitigate this, you can consider using the `undo irf mac-address persistent` command.

Logical Flow for IRF Split-Brain Detection and Recovery



[Click to download full resolution via product page](#)

Caption: IRF Split-Brain Detection and Recovery Workflow.

VXLAN Configuration

Question: My VXLAN tunnel is not coming up. How can I troubleshoot this?

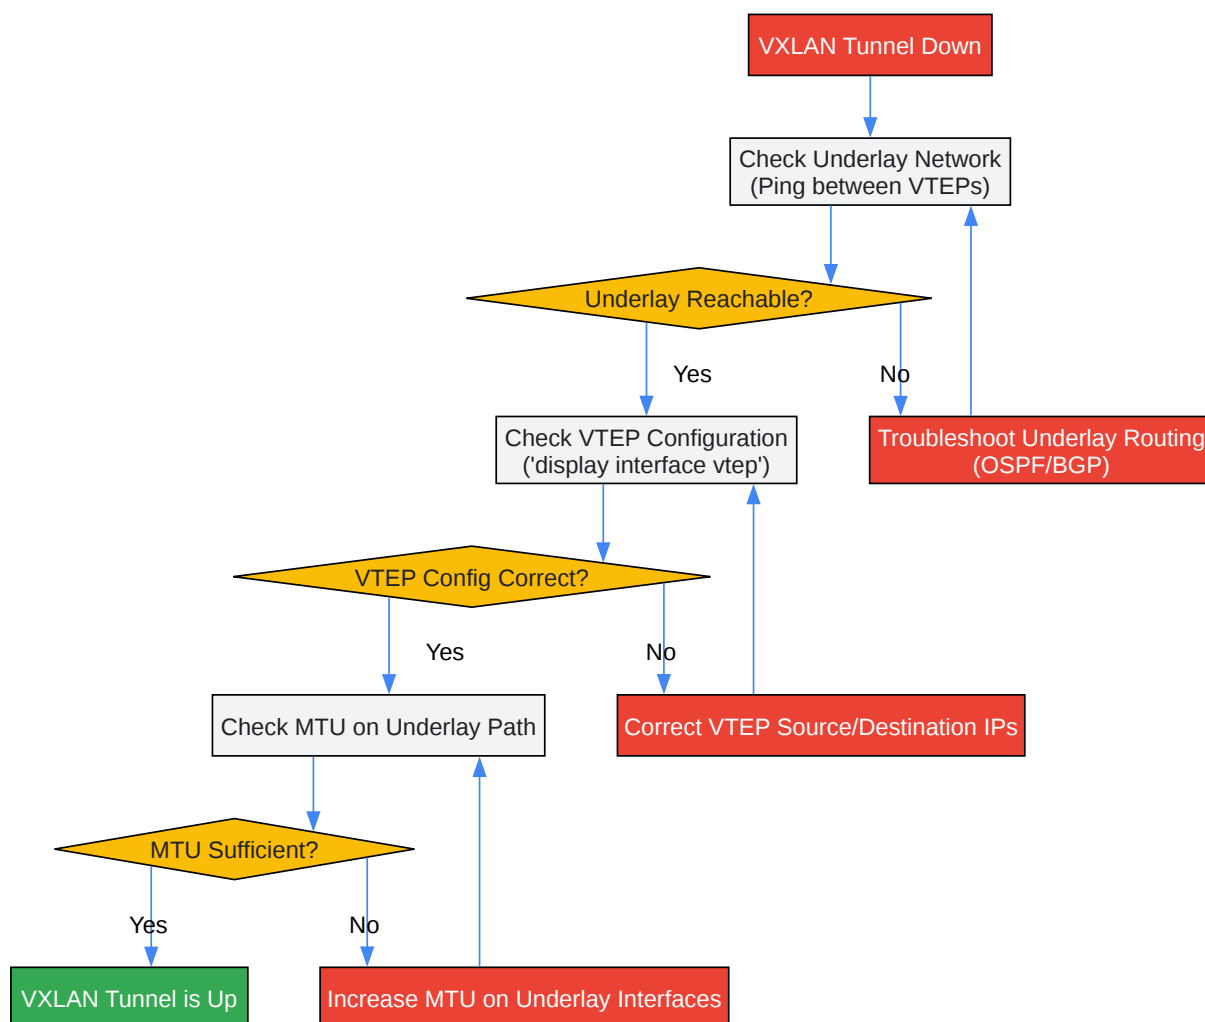
Answer:

VXLAN tunnel issues often stem from problems in the underlay network. Here's how to diagnose the problem:

Experimental Protocol: VXLAN Tunnel Troubleshooting

- **Verify Underlay Reachability:** The VTEPs (VXLAN Tunnel Endpoints) must be able to reach each other over the underlay network. Use the ping command with the source VTEP IP address to test connectivity to the destination VTEP IP address.
- **Check VTEP Configuration:**
 - Use the `display interface vtep` command to check the status of the VTEP interfaces.
 - Ensure that the source and destination IP addresses are correctly configured.
- **Inspect the Underlay Routing:** The routing protocol in the underlay network (e.g., OSPF or BGP) must be advertising the routes to the VTEP loopback interfaces.
- **MTU Mismatch:** VXLAN adds overhead to packets. Ensure that the MTU of all interfaces in the underlay network is large enough to accommodate the encapsulated packets.
- **Enable Overlay OAM:** For advanced troubleshooting, enable Overlay OAM (Operations, Administration, and Maintenance) to use tools like `ping vxlan` and `tracert vxlan` to test the VXLAN tunnel itself.

Troubleshooting Flow for VXLAN Tunnel Issues



[Click to download full resolution via product page](#)

Caption: A logical workflow for troubleshooting VXLAN tunnel connectivity.

BGP and OSPF Routing

Question: My BGP neighbor relationship is stuck in the "Idle" or "Active" state. What should I do?

Answer:

A BGP session stuck in "Idle" or "Active" state indicates a problem with establishing the TCP session between the peers.

Troubleshooting Protocol: BGP Neighbor States

- **Verify IP Connectivity:** Use the ping command to ensure there is basic IP reachability between the BGP neighbors. If you are using loopback interfaces for peering, make sure to source the ping from the loopback interface.
- **Check BGP Configuration:**
 - **Neighbor IP Address and AS Number:** Double-check that the peer as-number command is correctly configured on both routers.
 - **Update Source:** If using loopback interfaces, ensure the peer connect-interface LoopBack command is configured.
- **Check for Firewalls or ACLs:** Ensure that no firewalls or access control lists are blocking TCP port 179 between the BGP neighbors.
- **eBGP Multihop:** If the eBGP peers are not directly connected, you must configure peer ebgp-max-hop.

Question: My OSPF neighbors are stuck in a "2-Way" state. Is this an error?

Answer:

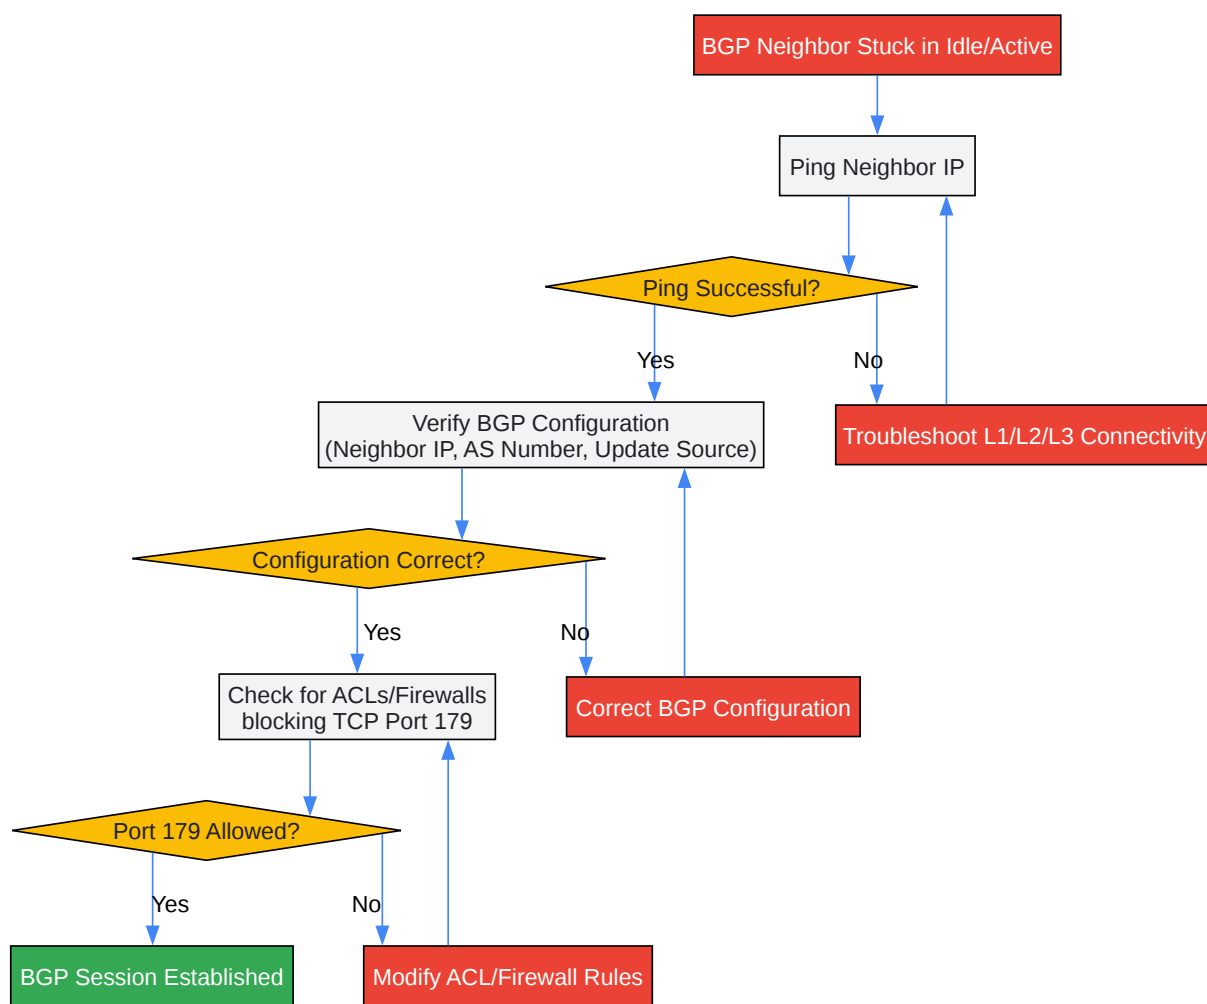
Not necessarily. On broadcast network types (like Ethernet), OSPF routers only form a full adjacency with the Designated Router (DR) and Backup Designated Router (BDR). The adjacency with other routers (DROthers) will remain in a "2-Way" state. This is normal OSPF behavior.

However, if you expect a full adjacency and it's stuck in "2-Way", it could indicate a configuration mismatch.

Troubleshooting Protocol: OSPF Adjacency Issues

- **Verify Area ID:** Both interfaces must be in the same OSPF area.
- **Check Timers:** The Hello and Dead timers must match on both sides of the link.
- **Authentication:** If using authentication, the key and authentication type must be identical.
- **Subnet Mask:** The interfaces must be on the same subnet with the same subnet mask.
- **Network Type:** Mismatched OSPF network types (e.g., one side configured as point-to-point and the other as broadcast) can prevent adjacency formation.

BGP Neighbor State Troubleshooting Flow



[Click to download full resolution via product page](#)

Caption: A step-by-step guide to troubleshooting BGP neighbor states.

FAQs

Q: What are some key performance metrics for the HPE 5945 switch series?

A: The HPE 5945 series offers high-performance switching for data center environments. Here are some of the key metrics:

Feature	Performance Metric
Switching Capacity	Up to 2.56 Tbps
Throughput	Up to 1904 MPPS (Million Packets Per Second)
40GbE Latency	Under 1μs
IRF Scalability	Up to 10 switches in a single fabric

Q: What are some useful display commands for troubleshooting?

A: Here is a table of commonly used diagnostic commands:

Command	Description
display irf	Shows the IRF topology and member status.
display vlan	Displays VLAN configuration and port membership.
display interface brief	Provides a summary of the status of all interfaces.
display ip routing-table	Shows the IP routing table.
display bgp peer	Displays the status of BGP peers.
display ospf peer	Shows the status of OSPF neighbors.
display vxlan	Provides information about VXLAN configurations.
display logbuffer	Shows the system log buffer.
display cpu-usage	Displays the CPU utilization of the switch.

Q: Are there any known bugs I should be aware of?

A: Yes, like any complex networking device, there can be software bugs. For example:

- A known issue exists in firmware release R6635 that can cause a slot to show up as "faulty" when forming an IRF stack with 100GbE interfaces. This is resolved in a later hotpatch (R6635H03) or subsequent general releases.
- A cosmetic bug in release R6555 may cause the display fan command to not show "FanDirectionFault" even when the fan airflow direction is mismatched with the preferred direction. This is fixed in release R6607.

It is always recommended to consult the latest release notes for your firmware version for a comprehensive list of known issues and resolved bugs.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. aboutnpnetworking.com [aboutnpnetworking.com]
- 2. arubanetworking.hpe.com [arubanetworking.hpe.com]
- 3. [scribd.com](https://www.scribd.com) [[scribd.com](https://www.scribd.com)]
- To cite this document: BenchChem. [HPE 5945 configuration errors and solutions]. BenchChem, [2025]. [Online PDF]. Available at: [<https://www.benchchem.com/product/b1669747#hpe-5945-configuration-errors-and-solutions>]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd

Ontario, CA 91761, United States

Phone: (601) 213-4426

Email: info@benchchem.com

