Enhanced Dual Stack

– Minimize IPv6 issues' impact on user experience during IPv6 introduction

XiPeng Xiao

Problem: Lack of Rule 1 Enforcement Leads to IPv6 Deployment Difficulty

RFC6724 Destination Address Selection:

- **Rule 1: Avoid unusable destinations**
  - **Comment: no method defined, generally not enforced**
- **Rule 2-5: … (not relevant)**
- **Rule 6: Prefer higher precedence**
  - **Comment: IPv6 preferred even when unusable**
- **Rule 7-10: … (not relevant)**

- **Deployers must ensure IPv6 works in all scenarios, or user experience may suffer**
  - **Comment: a lot of works as described in RFC7381 (Enterprise IPv6 deployment guidelines)**
- **Big IPv6 validation effort upfront, and some risks at switching on**



Enterprises with little IPv6 expertise & little need for large # of addresses have little incentive to deploy IPv6

If no action taken, many enterprises may remain in IPv4 for decades

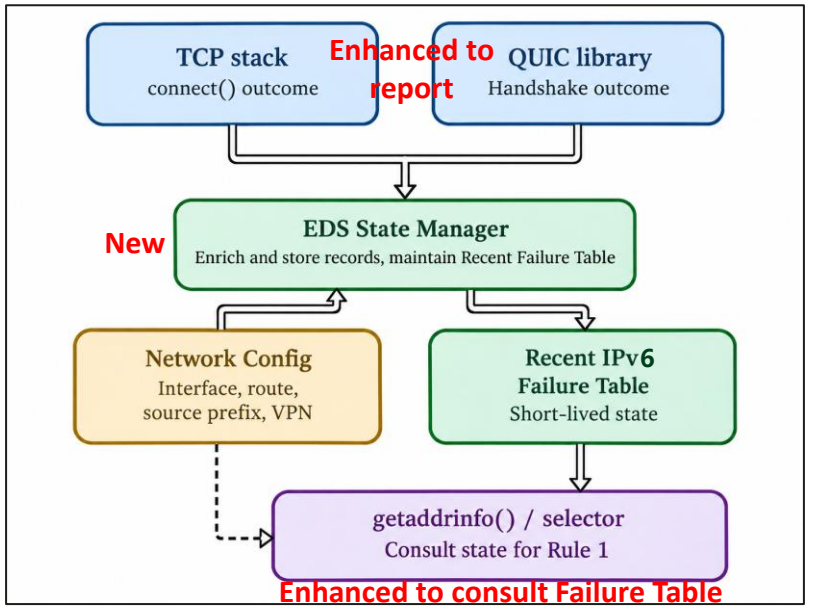
Proposed Solution: Find a Method to Identify “Unusable Destinations” for Rule 1

- **Cache TCP/QUIC connection failures so that Rule 1 can “Avoid unusable destinations”**
- **Make the failure information available for IPv6 problem diagnosis**

Recent failure table has specific connection info, failure type & timestamp

Field	Example
Network context	Corp-WiFi-A
IPv6 destination address	2001:db8:1234:5678::42
IPv6 source prefix or address	2001:db8:1111:1200::/64
Transport protocol and port	TCP/443
Failure type	timeout
Timestamp	2026-09-17T14:32:18+02:00

Mechanisms to construct the failure table, only State Manager is new, others are enhancement to existing mechanisms



Changes from Previous Version: No Longer Depends on Performance Measurement or Happy Eyeballs

- **Performance-based destination selection → connection-result-based.**
  - **No more performance measurements (e.g. RTT, loss, etc) → simplified**
- **Removed the dependency on Happy Eyeballs (HE) – HE & EDS are different ways to "Avoid unusable destinations"**

Aspect	HE	EDS
Approach	Race IPv6/IPv4 connections; use the winner	Cache recent failures; avoid failed destinations
Beneficiaries	Applications implementing HE	All applications using EDS-optimized OS/platform
First IPv6 failure	Minimizes delay through racing	May time out, as with current RFC 6724
Complexity	More complex; requires racing	Simpler; caches connection failures

- **Changed the incremental deployment approach**

Incremental Deployment Method: Can Start with IPv6 Pioneer Hosts Supporting EDS

- **Start small: Divide hosts into IPv6-pioneers and regular hosts**
  - **Pioneer hosts: EDS-capable**
  - **Regular hosts: remain IPv4 at the beginning**
- **Practical validation: Perform IPv6 validation as much as practical using RFC7381 guidelines**
- **Enable IPv6: on pioneer hosts and supporting infrastructure (e.g., networks)**
- **Diagnose and fix: Use pioneer-host diagnostics to identify and correct IPv6 problems**
- **Expand gradually: Enable IPv6 on more hosts, applications, services, and networks as confidence grows; eventually migrate toward IPv6-mostly/IPv6-only.**

Benefit: Turn IPv6 Rollout from Climbing a Cliff into Climbing a Slope, Suitable for Enterprises with Little IPv6 Expertise

Now: do lots of IPv6 validation works upfront



EDS: roll out IPv6 incrementally

5. Gain sufficient IPv6 confidence, migrate to IPv6-only / mostly
4. Fix IPv6 problems, grow IPv6 traffic%
3. Enable IPv6, use EDS stats to help diagnose IPv6 problems
2. Do as much IPv6 validation as practical
1. Upgrade pioneer hosts to support EDS

The EDS tradeoff: ask a few OS companies to optimize host behavior so that many enterprises can introduce IPv6 easier

- **Request to WG:**
  - **Please review and comment**
  - **Adopt the draft?**