

# BGP技术：路由反射器

Ender.joe(周亚军)  
RS CCIE,SP CCIE,思科认证讲师#34708  
RS & SP CCIE讲师, Yeslab (上海)

# Introducing Route Reflectors



# Outline

## Overview

IBGP Scalability Issues in a Transit AS-RR在一个大型传输网络的作用

Route Reflector Split-Horizon Rules-RR的水平分割原则

Redundant Route Reflectors-冗余的RR

Route Reflector Clusters-RR簇


Additional Route Reflector Loop-Prevention Mechanisms-RR的防环机制

Summary

# IBGP Scalability Issues in a Transit AS

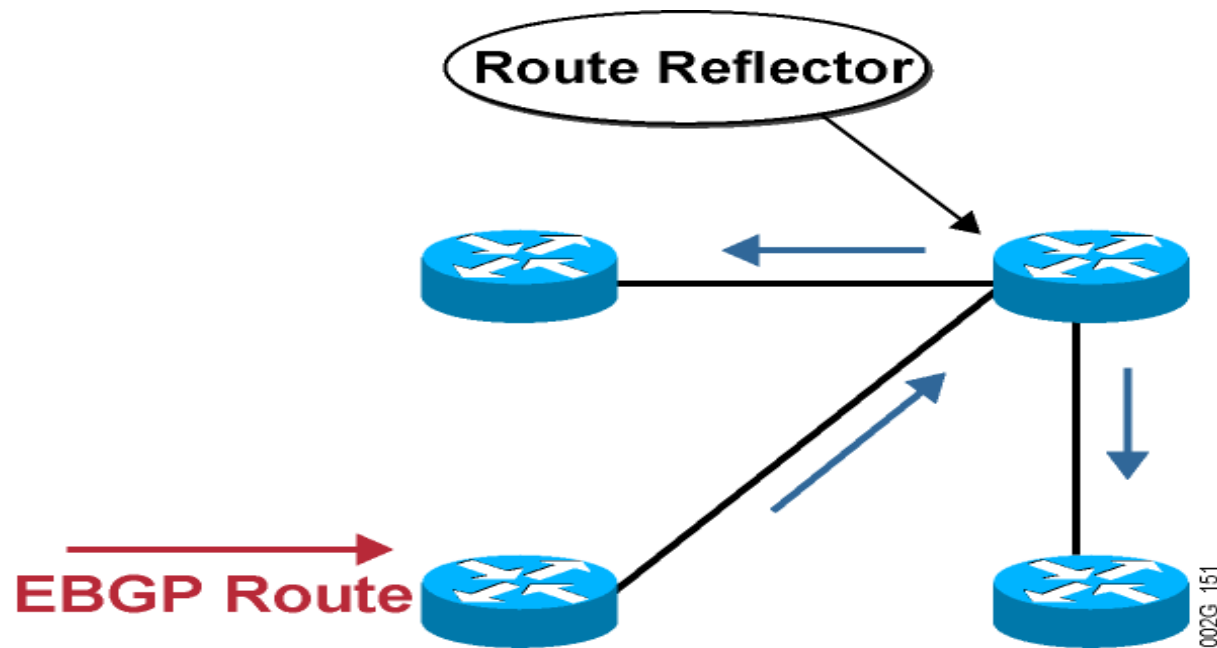
- IBGP requires a full mesh between all BGP-speaking routers.
  - Large number of TCP sessions
  - Unnecessary duplicate routing traffic
- Solutions
  - Route reflectors modify IBGP split-horizon rules (RR修改了IBGP的水平分割原则)
  - BGP confederations modify IBGP AS-path processing (BGP联邦修改了IBGP的AS-path)

# Route Reflector Split-Horizon Rules



Classic IBGP: IBGP routes are not propagated to other IBGP peers.

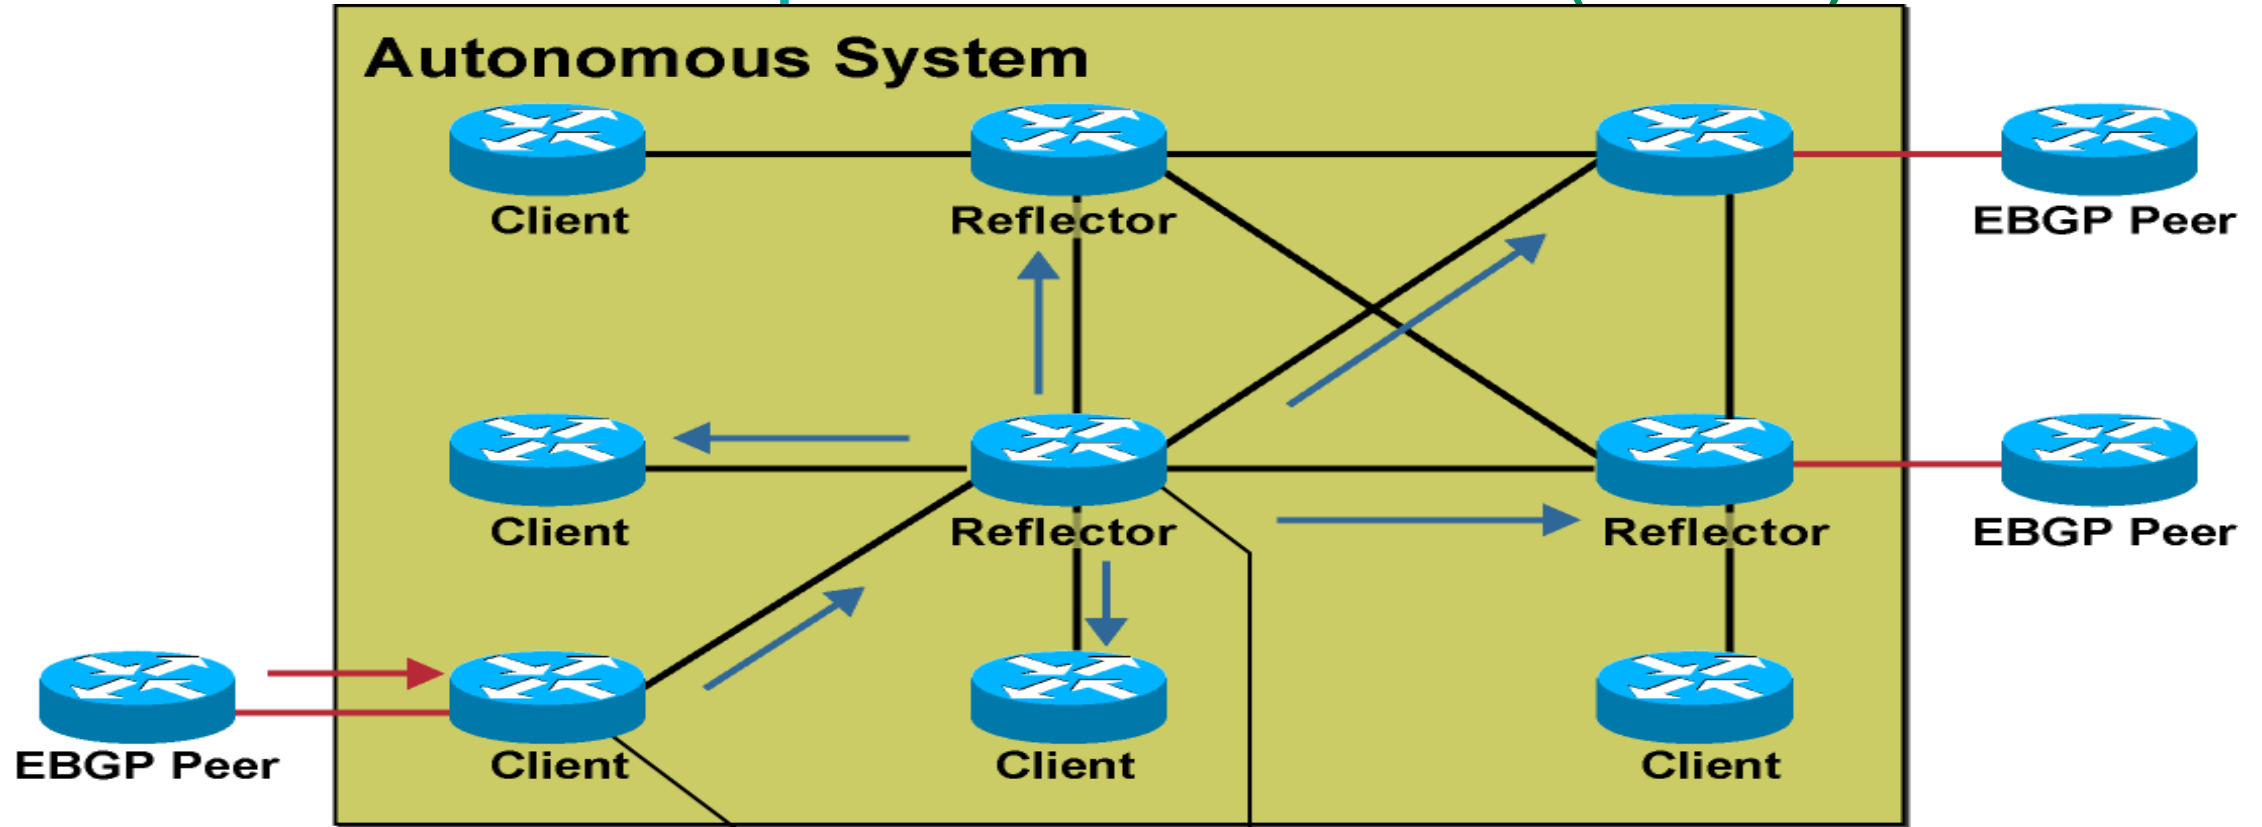
Full mesh of IBGP peers is therefore required.



Route reflector can propagate IBGP routes to other IBGP peers (RR通告IBGP路由给其他IBGP peer)

Full mesh of IBGP peers is no longer required (全互联不再需要)

# Route Reflector Split-Horizon Rules (Cont.)



2、从一个RR的客户端收到的路由通告给其他所有邻居  
(包括客户端和非客户端)

1、从EBGP收到的通告给所有的内部邻居

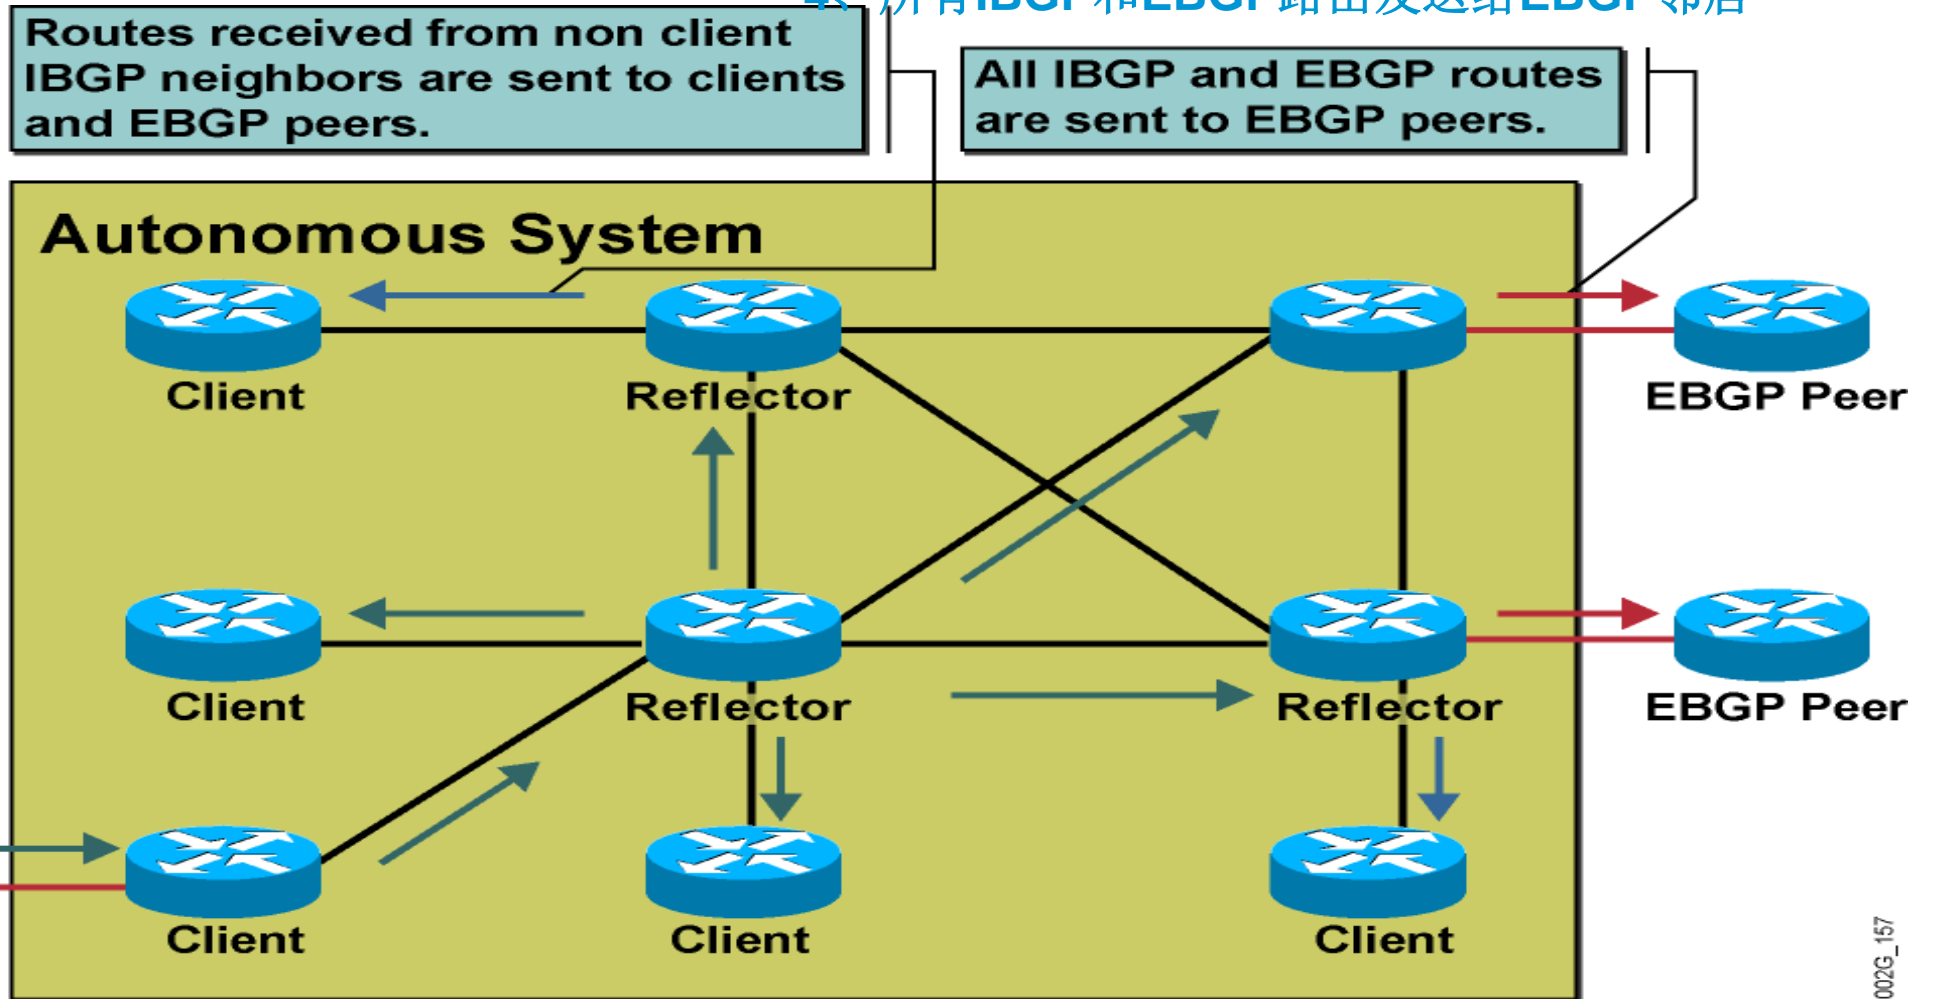
Routes received from a client are propagated to all other peers.

Routes received from external peers are propagated to all internal peers.

# Route Reflector Split-Horizon Rules (Cont.)

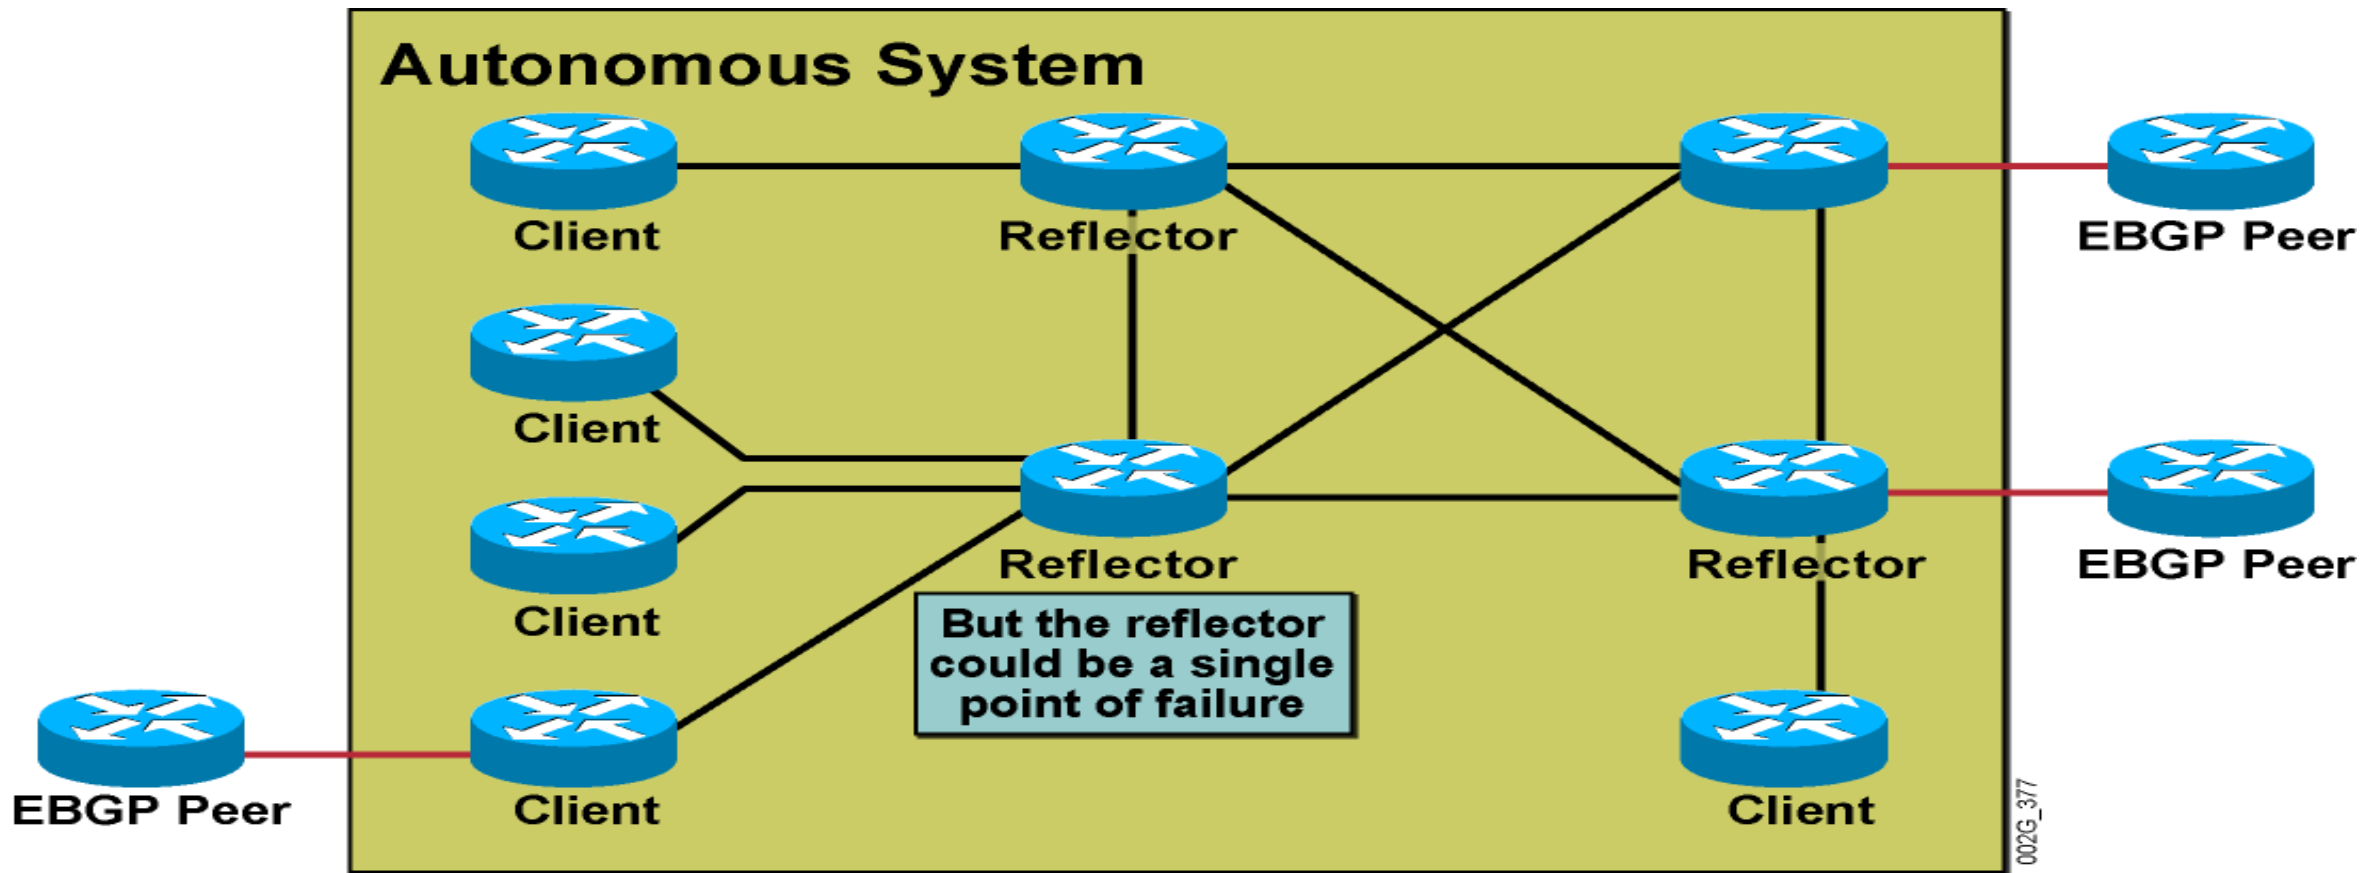
4、所有IBGP和EBGP路由发送给EBGP邻居

3、从一个非客户端的IBGP邻居收到的路由会发送给客户端和EBGP邻居



002G\_157

# Redundant Route Reflectors



**Design Requirement: Route reflectors must be redundant.**

RR必须设计成为冗余的，否则会成为单点故障点

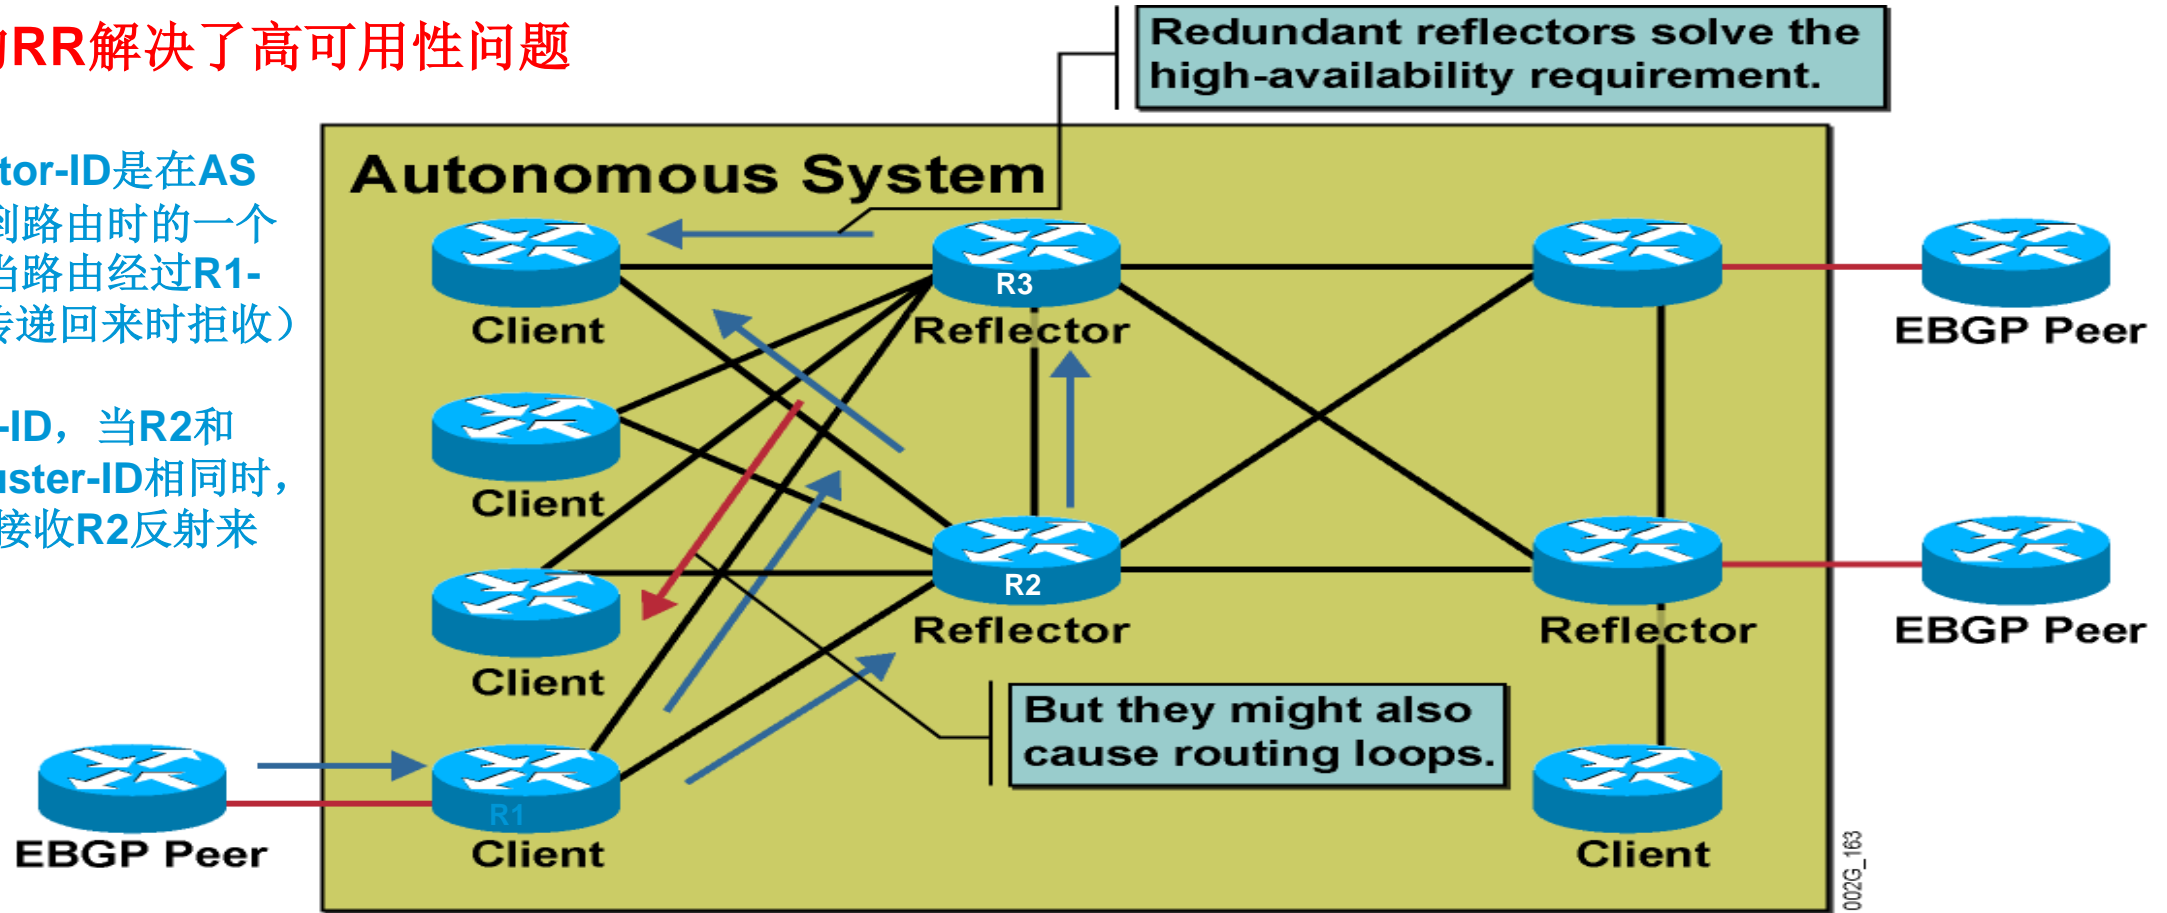


# Redundant Route Reflectors (Cont.)

冗余的RR解决了高可用性问题

Originator-ID是在AS边界得到路由时的一个标识（当路由经过R1-R2-R3传递回来时拒收）

Cluster-ID，当R2和R3的cluster-ID相同时，R3不再接收R2反射来的路由



The concept of "clusters" is introduced to prevent IBGP routing loops between route reflectors.

簇的概念用于在RR之间阻止环路

# Route Reflector Clusters

A group of redundant route reflectors and their clients form a cluster-冗余的RR和他们的客户端形成一个簇

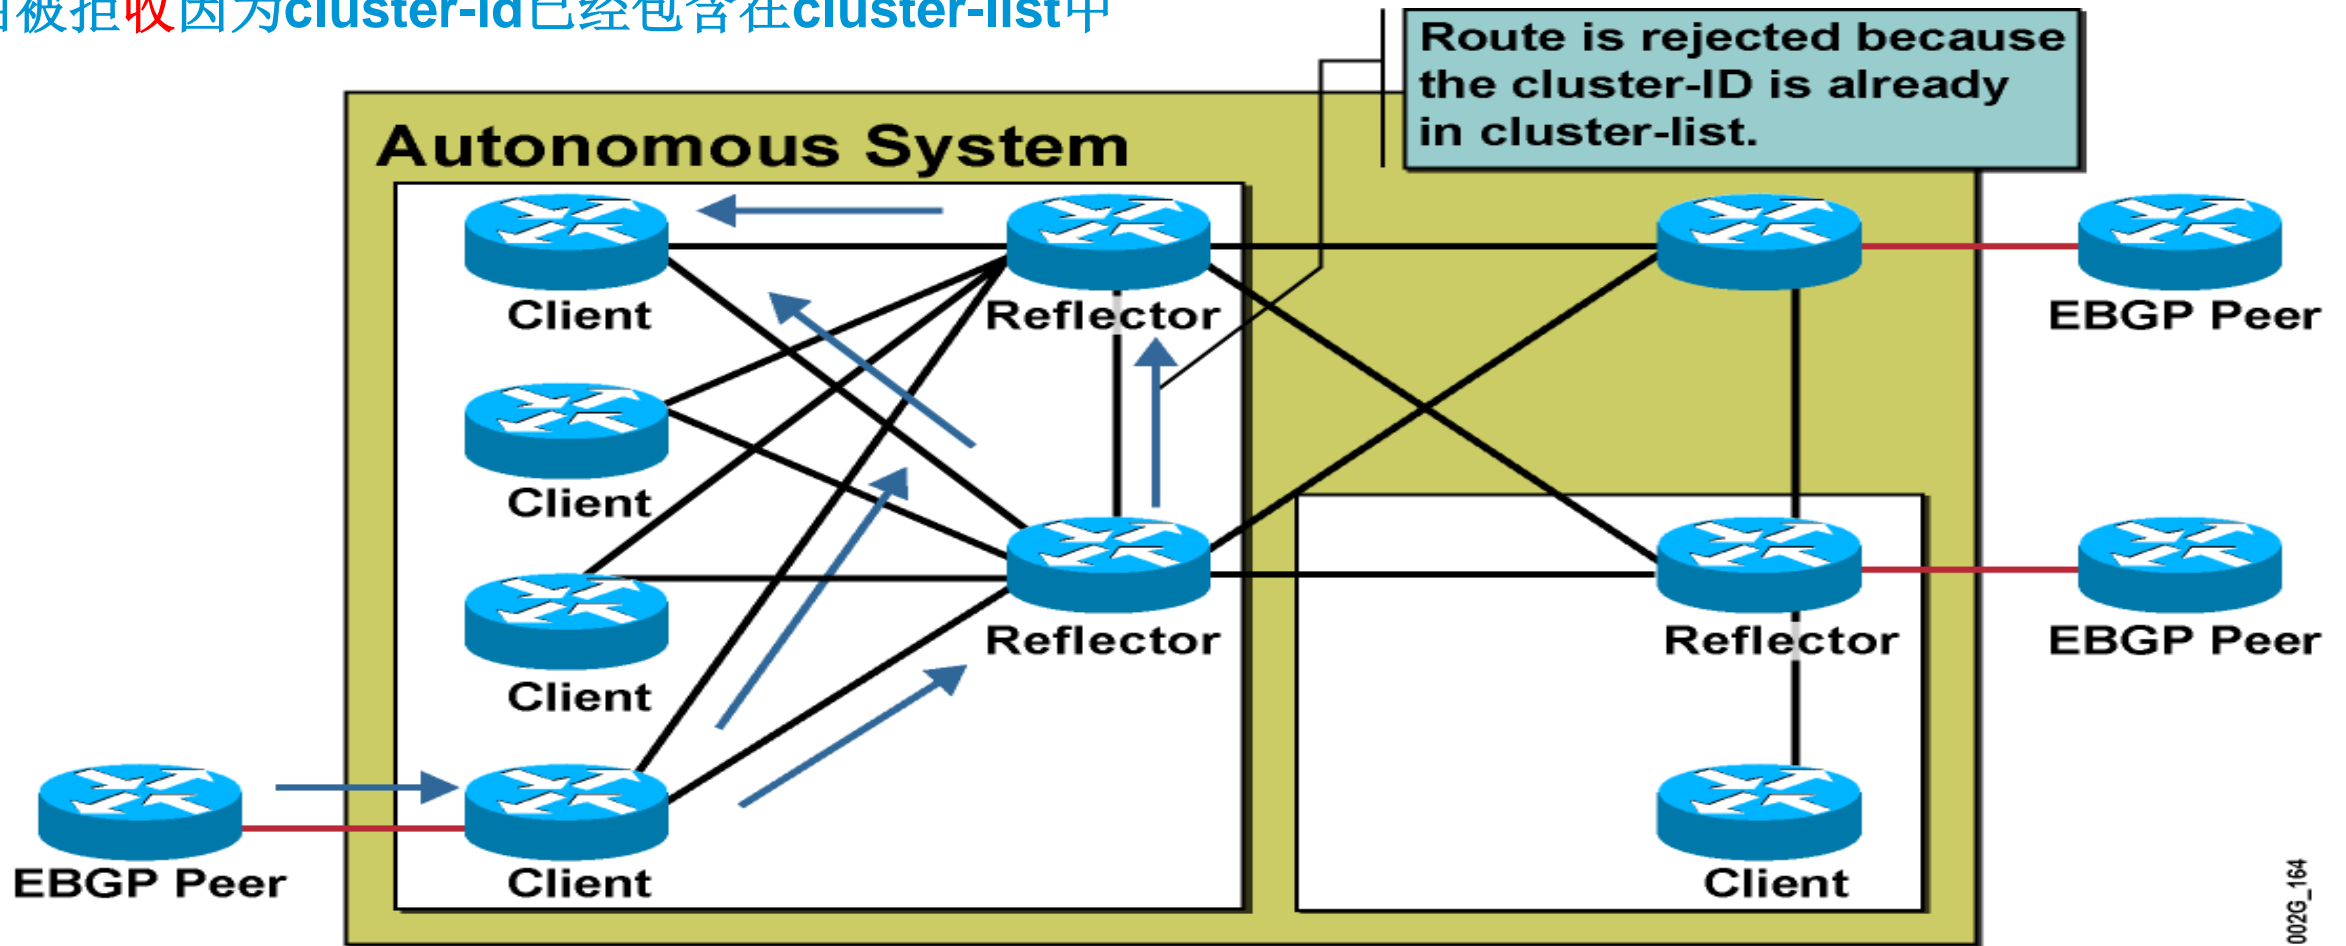
Each cluster must have a unique cluster-ID-每个簇有一个单独的cluster-ID

Each time a route is reflected, the cluster-ID is added to the cluster-list BGP attribute-当路由被反射时，会在BGP的cluster-list属性中增加cluster-ID

The route that already contains the local cluster-ID in the cluster-list is not reflected-收到的路由已经包含了本身的cluster-ID将不会接收

# Route Reflector Clusters (Cont.)

路由被拒收因为cluster-id已经包含在cluster-list中



0026\_164

# Additional Route Reflector Loop-Prevention Mechanisms

Every time a route is reflected, the router-ID of the originating IBGP router is stored in the originator-ID BGP attribute. 一条路由被反射时，起源的iBGP路由器的RID会被收集为originator-ID属性

A router receiving an IBGP route with originator-ID set to its own router-ID ignores that route. 当收到一个originator-ID和本路由器的RID一致时忽略该路由

The BGP path selection procedure is modified to take into account cluster-list and originator-ID-第11条和第12条选路原则被修改

# Configuring Route Reflectors

Configure cluster-ID on route reflectors.

Configure BGP neighbors as route reflector clients on the route reflectors.

No configuration is needed on the route reflector clients.

# Configuring Route Reflectors (Cont.)

```
router(config-router)#
```

```
bgp cluster-id cluster-id
```

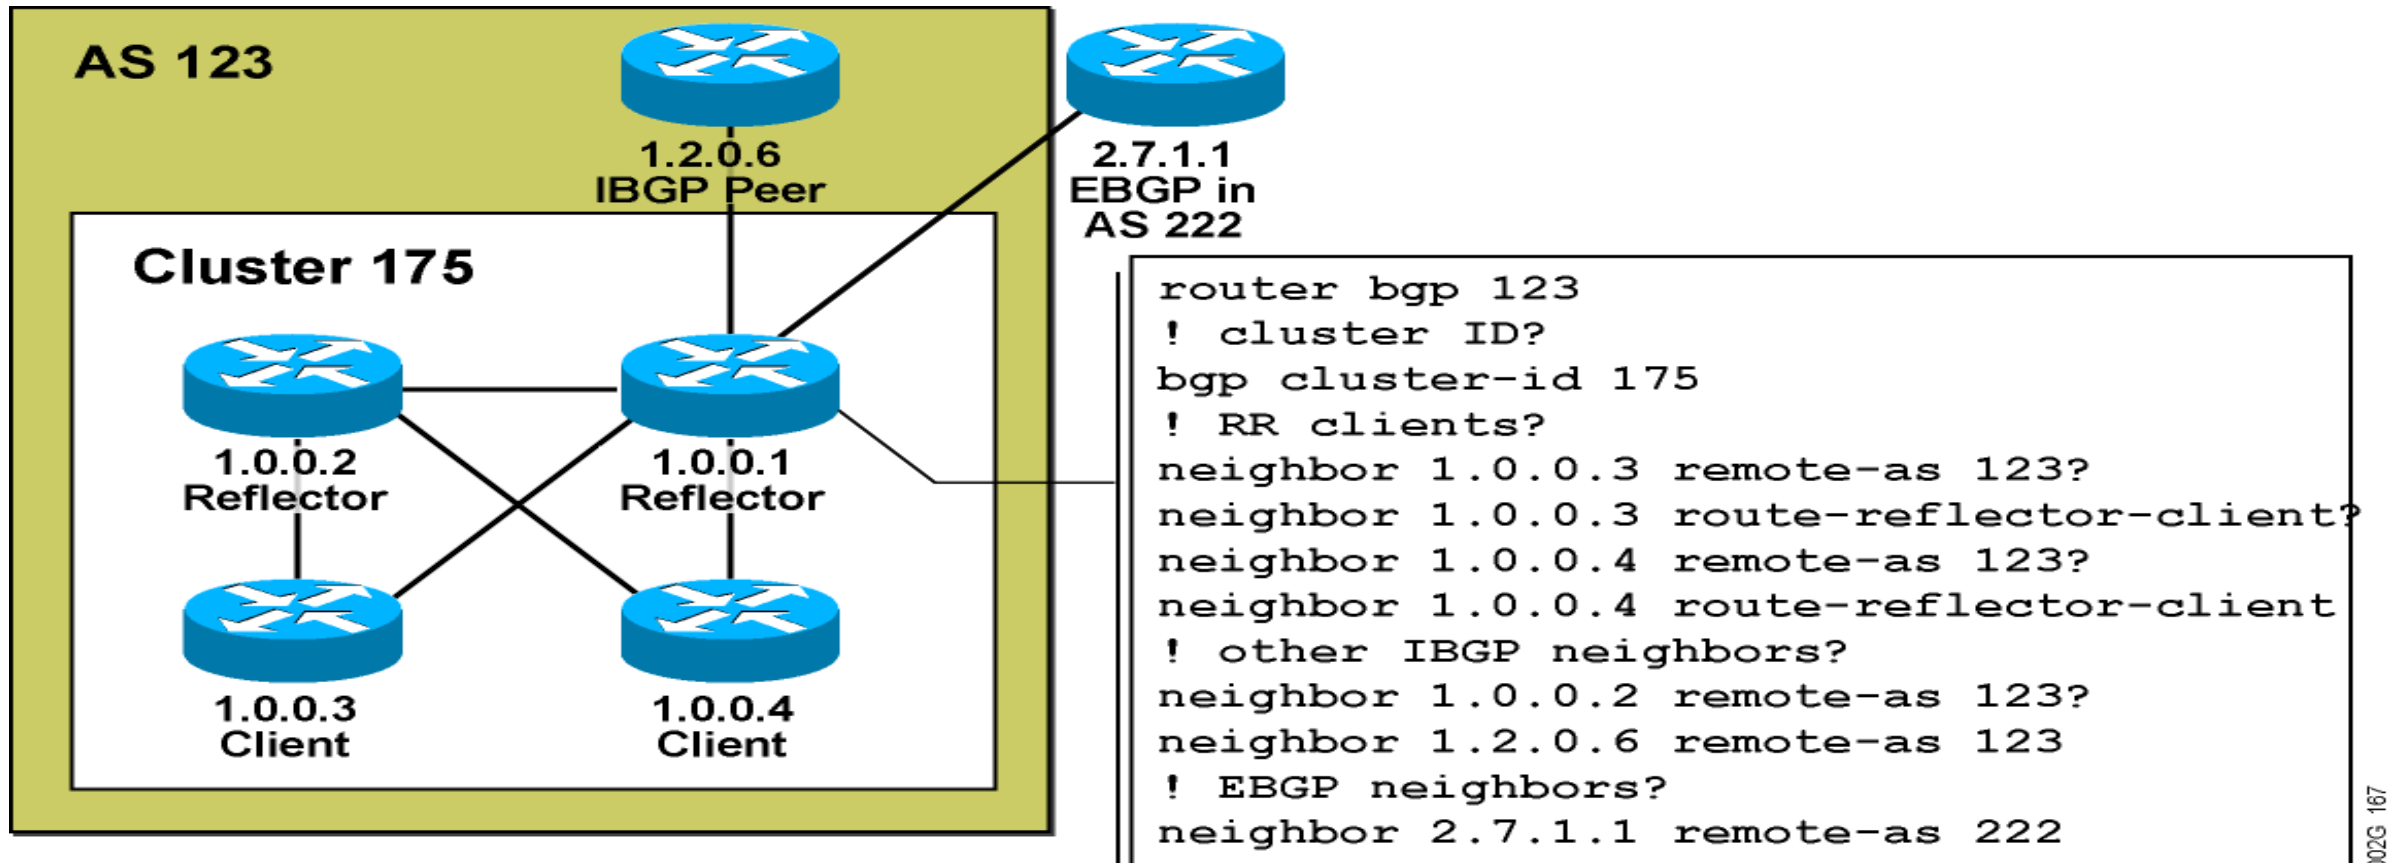
- **Optionally assigns a cluster-ID to the route reflector (default value is router-ID)**
- **Required only for clusters with redundant reflectors**
- **Cluster-ID cannot be changed after the first client is configured**

```
router(config-router)#
```

```
neighbor ip-address route-reflector-client
```

- **Configures an IBGP neighbor to be a client of this reflector**

# Configuring Route Reflectors (Cont.)



002G\_167

# Monitoring Route Reflectors

```
router#
```

```
show ip bgp neighbors
```

- **Displays whether a neighbor is a route reflector client**

```
router#
```

```
show ip bgp network [mask]
```

- **Displays additional path attributes (originator-ID and cluster-list)**



# Monitoring Route Reflectors (Cont.)

```
router# show ip bgp neighbors 1.0.0.1
BGP neighbor is 1.0.0.1, remote AS 213, internal link
Index 1, Offset 0, Mask 0x2
Route-Reflector Client
BGP version 4, remote router ID 11.0.0.1
BGP state = Established, table version = 5, up for 01:33:24
Last read 00:00:24, hold time is 180, keepalive interval is
60 seconds
Minimum time between advertisement runs is 5 seconds
Received 257 messages, 0 notifications, 0 in queue
Sent 264 messages, 0 notifications, 0 in queue
Connections established 5; dropped 4
Last reset 01:33:33, due to : User reset request
No. of prefix received 1
```

002G\_168

# Monitoring Route Reflectors (Cont.)

```
rtr-a# show ip bgp 11.0.0.0
BGP routing table entry for 11.0.0.0/8, version 3
Paths: (1 available, best #1, advertised over IBGP)
  Local, (Received from a RR-client)
    1.0.0.1 (metric 40640000) from 1.0.0.1 (11.0.0.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Routes received from the client as seen on the reflector

```
rtr-b# sh ip bgp 14.0.0.0
BGP routing table entry for 14.0.0.0/8, version 30
Paths: (1 available, best #1)
  Not advertised to any peer
  Local
    1.0.0.3 (metric 41152000) from 1.0.0.2 (14.1.2.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 14.1.2.3, Cluster list: 0.0.2.55
```

Reflected routes as seen on the client

# Summary

BGP route reflectors were introduced to free the network designers from IBGP full-mesh requirements that prevent large networks from scaling.

BGP route reflectors modify IBGP split-horizon rules in that all routes that are received from a route reflector client are sent to all other IBGP neighbors, and all routes that are received from a nonclient IBGP neighbor are sent to all route reflector clients.

A route reflector is a single point of failure, and therefore redundancy should be implemented in a network containing route reflectors.

# Summary (Cont.)

Route reflector clusters were introduced in the BGP route reflector architecture to support redundancy, preventing IBGP routing loops in redundant route reflector designs.

The originator-ID and cluster-list BGP attributes were introduced to prevent routing loops in route reflector environments.

Thank you.

