

Analysis on Two Methods in Ingress Local Protection

Contents

- Introduction
- Descriptions on Two Methods
- Analysis on Two Methods

Introduction

Important to have **Ingress/Egress Local Protection**

- **Faster** than E2E global protection for ingress/egress
- **More Scalable** since keeping less states and using less resources
- **Easier** to operate and maintain

In Ingress Local Protection:

Relay-Message or **Proxy-Ingress** can be used to transfer information about ingress local protection from primary ingress to backup ingress

They are **described** in the following slides

Detailed analysis on them are given for your reference

Contents

- Introduction
- Descriptions on Two Methods
 - Changes for Relay-Message vs. Proxy-Ingress
 - Relay-Message with Example
 - Proxy-Ingress with Example
- Analysis on Two Methods

Changes for Relay-Message vs. Proxy-Ingress

Changes for Relay-Message

- 1) Primary ingress sends Path messages with **Ingress-Protection** object to backup ingress after the primary LSP is set up.
- 2) Backup ingress creates backup LSP to locally protect the primary ingress after receiving Path message with **Ingress-Protection** object, and sends Resv message with Ingress-Protection object to primary ingress.
- 3) Primary ingress records the status of ingress protection after receiving Resv message with **Ingress-Protection** object.

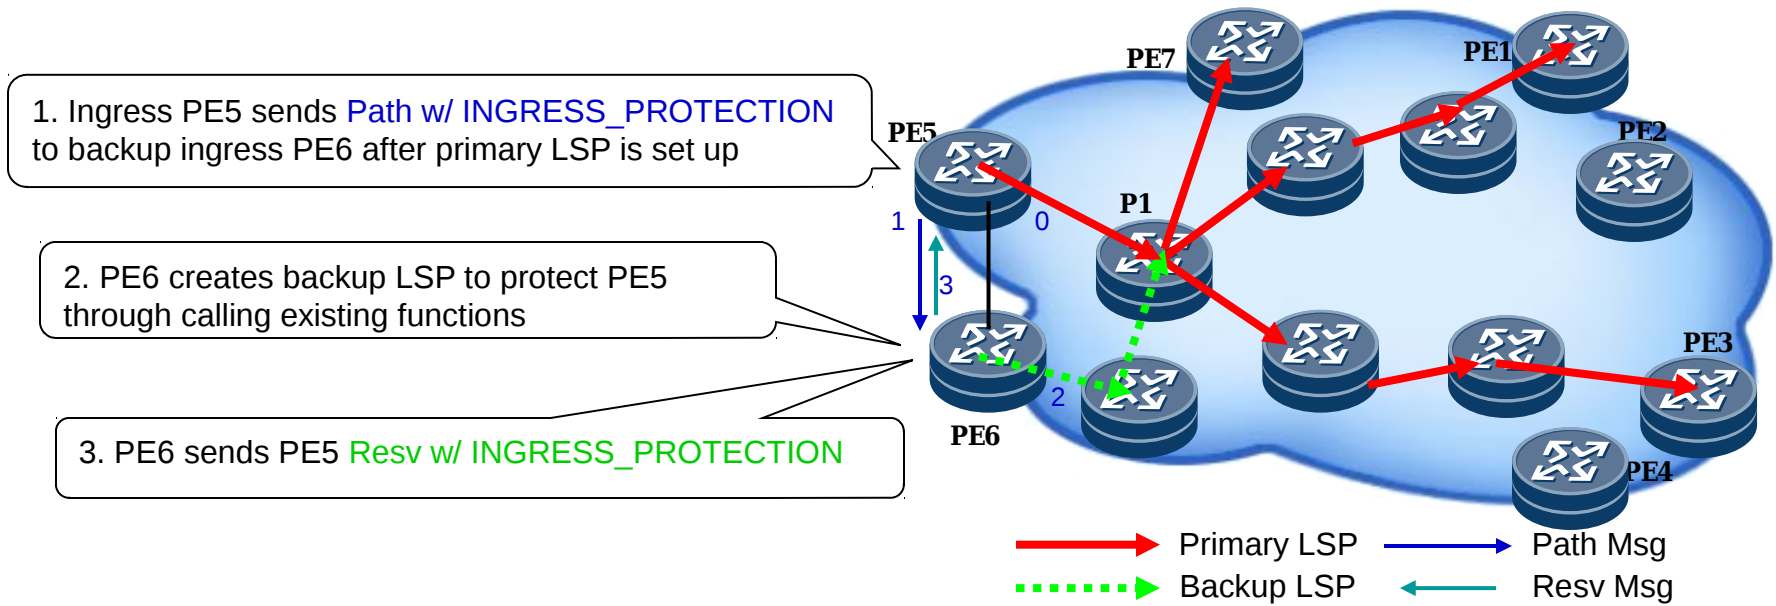
Changes for Proxy-Ingress

- 1) Primary ingress handles the configuration of proxy-ingress or generates the information for the proxy-ingress and makes sure that the proxy-ingress address generated does not cause a loop.
- 2) Primary ingress specially processes all possible abnormal cases happening in the backup ingress and in the path segment between the proxy ingress (i.e., the primary ingress), backup ingress and the primary ingress. These are changes to the existing RSVP-TE protocol, especially mixed with signaling for the primary LSP.
- 3) Primary ingress changes the path for the primary LSP. The new path for the LSP will be: the proxy-ingress (i.e., the primary ingress), the backup ingress, the primary ingress, the next hop(s) of the primary ingress, and so on.
- 4) Primary ingress adds a new object (**Ingress-Protection**) into the Path and Resv messages for the primary LSP to the backup ingress.
- 5) Primary ingress specially handles the Path and Resv messages w/ **Ingress-Protection** for the primary LSP from and/or to the backup ingress.
- 6) Backup ingress specially handles the Path and Resv messages w/ **Ingress-Protection** for the primary LSP from and/or to the primary ingress. The procedures on the backup ingress for specially handling the Path and Resv messages are different from those on the primary ingress.

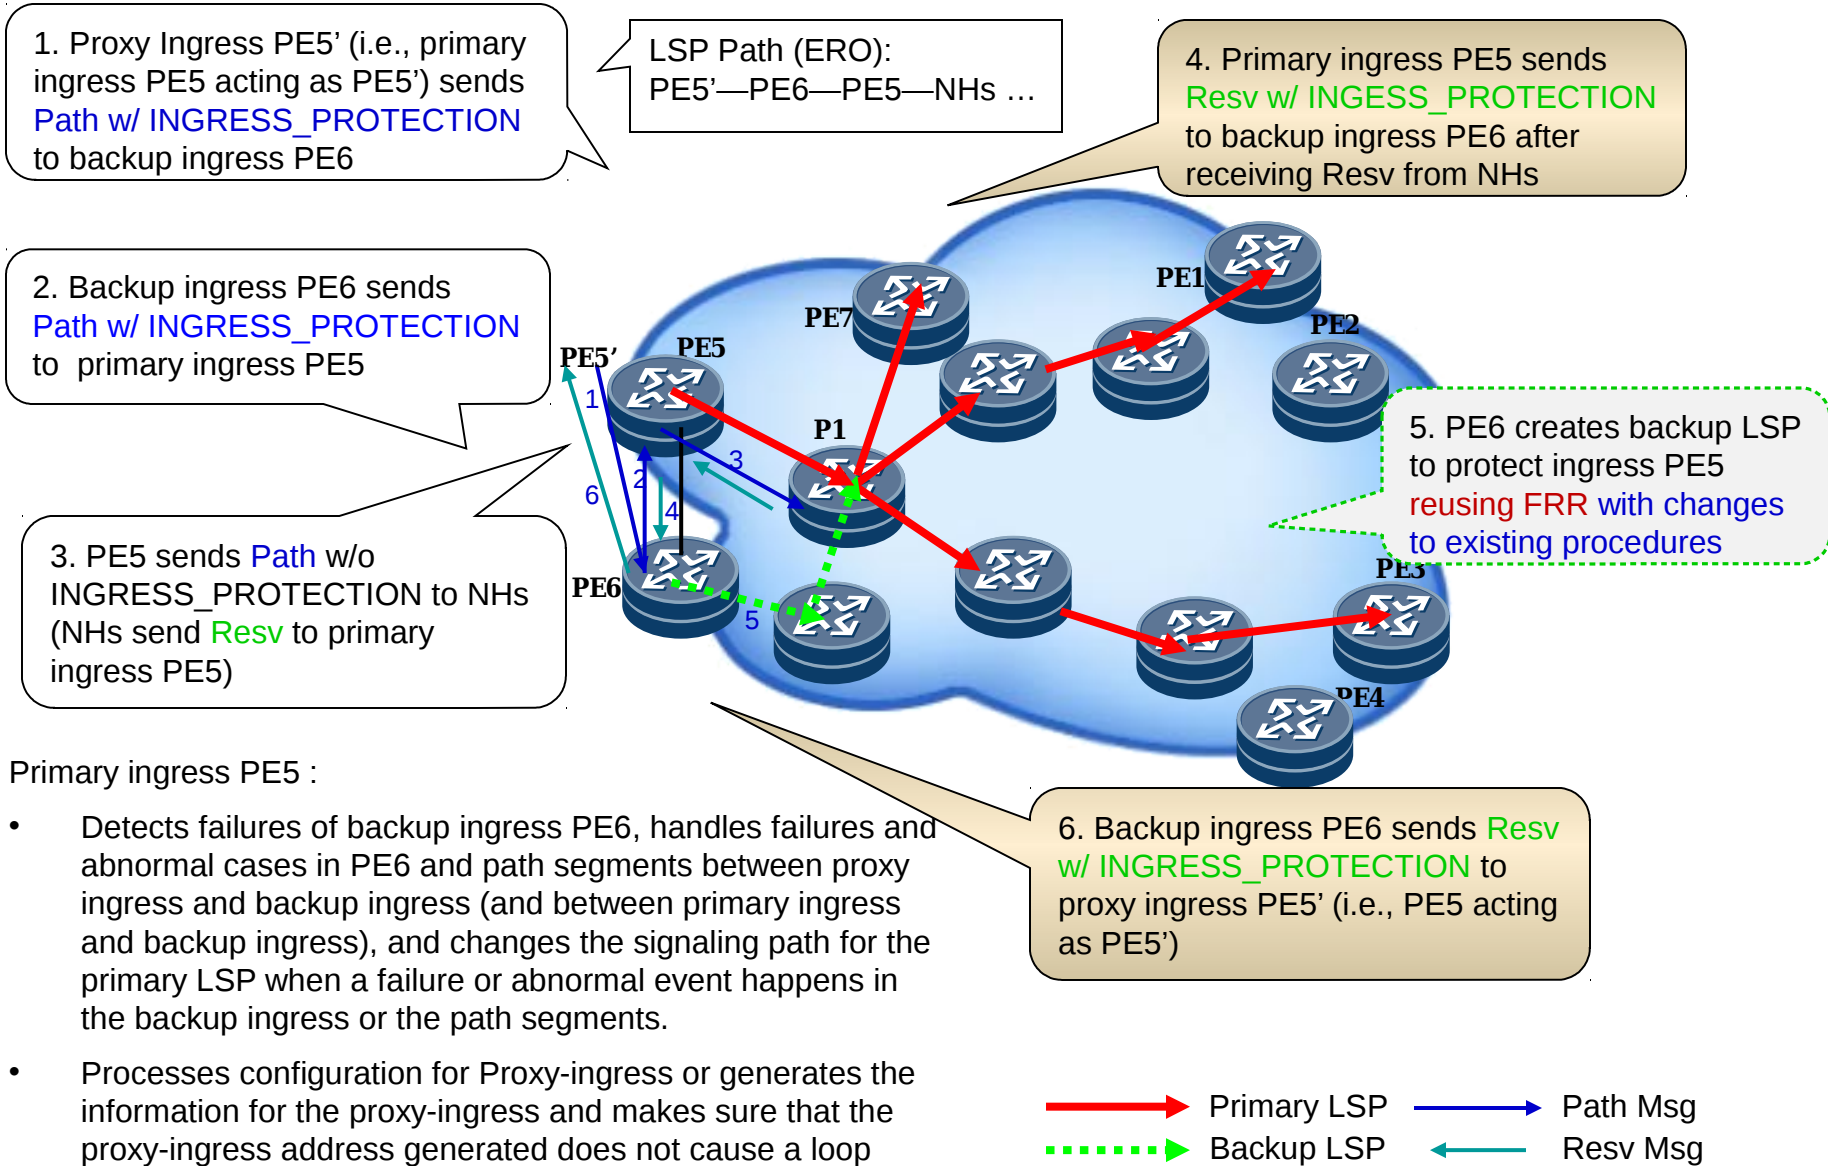
When the primary ingress fails, the backup ingress can not get any Path messages from the primary ingress (i.e., the proxy-ingress or the primary ingress), thus it must keep the Path message(s) originally received from the primary ingress, update the message(s) and put the message(s) into the bypass LSP tunnel to the next hop(s) of the primary ingress.

When the primary ingress fails, the backup ingress can not send any Resv message(s) to its primary ingress (i.e., its previous hop proxy-ingress or the primary ingress), thus it should keep the Resv message(s) originally received and update the message(s) such as setting Protection-in-use.

Relay-Message Method with Example



Proxy-Ingress Method with Example



Contents

- Introduction
- Descriptions on Two Methods
- Analysis on Two Methods
 - Configurations
 - Primary LSP Dependency
 - Message Overhead
 - Special Handlings on Primary Ingress
 - Special Handlings on Backup Ingress
 - Backup LSP Creation
 - Primary LSP Setup Time
 - Session Maintenance
 - Scalability
 - Summary on Analysis

Configurations

Relay-Message Method:

- Configure backup ingress

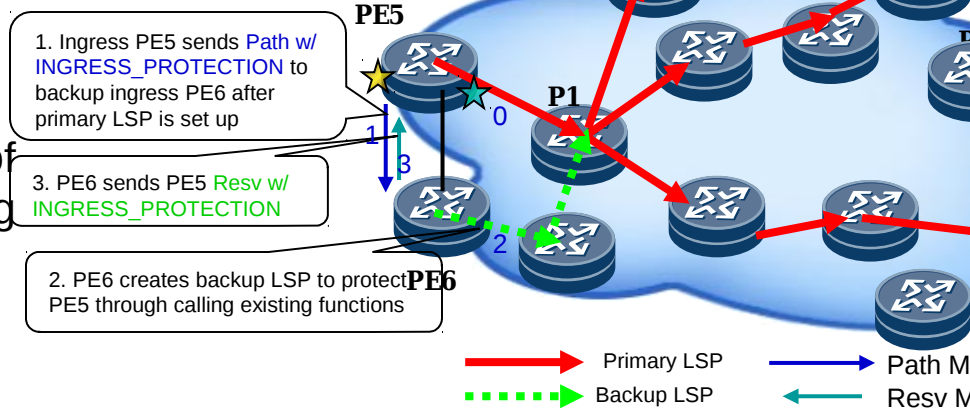
Proxy-Ingress Method:

- Configure backup ingress
- **Configure proxy-ingress** or generate the information for the proxy-ingress and make sure that the proxy-ingress address generated does not cause a loop

Primary LSP Dependency

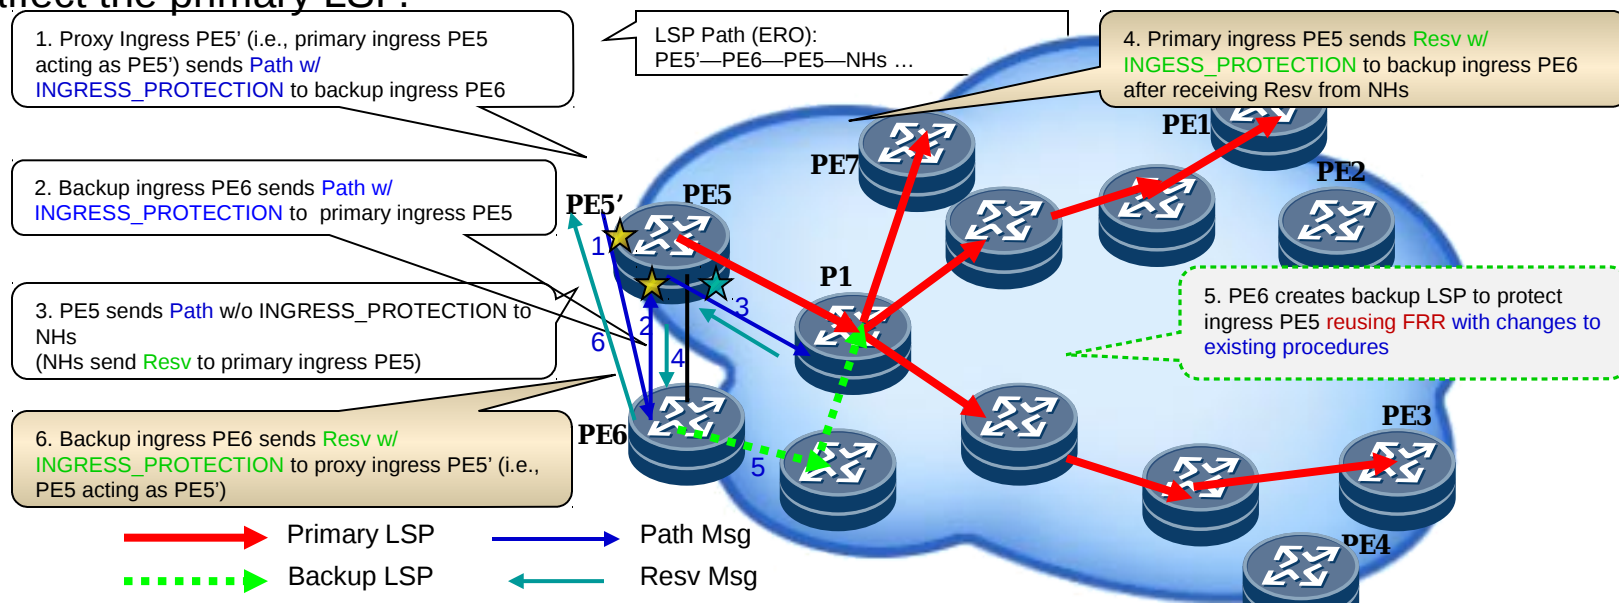
Relay-Message Method:

- Primary LSP is independent of the backup ingress. The establishment of primary LSP is not touched by adding ingress protection



Proxy-Ingress Method:

- The primary LSP depends on the backup ingress somehow. The creation of primary LSP is changed for providing ingress protection. For example, the signaling messages for the primary LSP goes from the primary ingress (acting as the proxy ingress) to the backup ingress, to the primary ingress and then to the next hop(s) of the primary ingress. Thus abnormal events on backup ingress may affect the primary LSP.




Message Overhead

Relay-Message Method (2 messages):

- Path Message with INGRESS_PROTECTION is sent
 - from primary ingress to backup ingress
- Resv Message with INGRESS_PROTECTION is sent
 - from backup ingress to primary ingress

1. Ingress PE5 sends Path w/ INGRESS_PROTECTION to backup ingress PE6 after primary LSP is set up

3. PE6 sends PE5 Resv w/ INGRESS_PROTECTION



Proxy-Ingress Method (4 messages):

- Path Message with INGRESS_PROTECTION is sent
 - from proxy-ingress (i.e., primary ingress) to backup ingress and
 - from backup ingress to primary ingress
- Resv Message with INGRESS_PROTECTION is sent
 - from primary ingress to backup ingress and
 - from backup ingress to proxy-ingress (i.e., primary ingress)


1. Proxy Ingress PE5' (i.e., primary ingress PE5 acting as PE5') sends Path w/ INGRESS_PROTECTION to backup ingress PE6

2. Backup ingress PE6 sends Path w/ INGRESS_PROTECTION to primary ingress PE5

LSP Path (ERO): PE5'—PE6—PE5—NHs ...

4. Primary ingress PE5 sends Resv w/ INGRESS_PROTECTION to backup ingress PE6 after receiving Resv from NHs

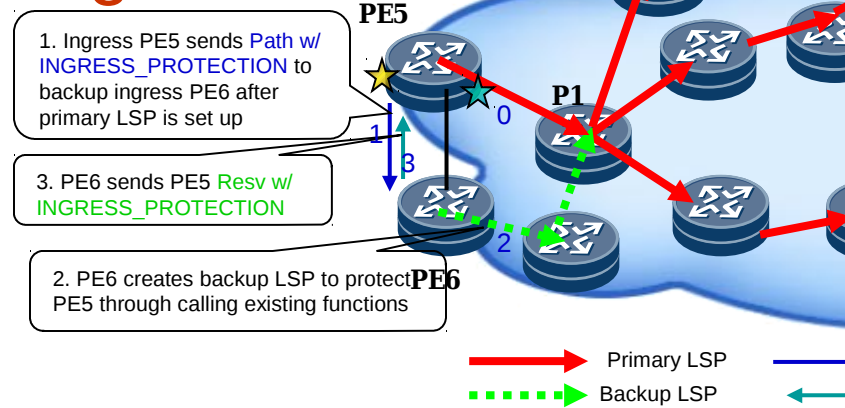
6. Backup ingress PE6 sends Resv w/ INGRESS_PROTECTION to proxy ingress PE5' (i.e., PE5 acting as PE5')



Special Handlings on Primary Ingress


Relay-Message Method:

1. sends **Path with INGRESS_PROTECTION** to backup ingress after primary LSP is set up
2. stores states after receiving **Resv with INGRESS_PROTECTION** from backup ingress
3. removes session to backup ingress when tearing down primary LSP (or failure or abnormal events in backup ingress occur)
4. refreshes session to backup ingress when refreshing primary LSP



Proxy-Ingress Method:

1. changes the ERO for primary LSP (ERO: proxy-ingress → backup ingress → primary ingress → next hop → ...)
2. sends **Path with INGRESS_PROTECTION** to backup ingress after changing the ERO
3. sends **Path without INGRESS_PROTECTION** to next hop after receiving **Path with INGRESS_PROTECTION** from backup ingress
4. After receiving Resv without INGRESS_PROTECTION from next hop, it does not assign a normal label for the interface from its upstream node (backup ingress), or write a cross connect from its upstream node to its next hop; it sends **Resv with INGRESS_PROTECTION** to its upstream node (i.e., backup ingress).
5. After receiving **Resv with INGRESS_PROTECTION** from backup ingress, stores states and write LFIB for primary LSP
6. detects failures of the backup ingress, handles failures and other abnormal events happening in the backup ingress and the path segments between proxy ingress and backup ingress (and between primary ingress and backup ingress), and changes the signaling path for the primary LSP when a failure or abnormal event happens in the backup ingress or the path segments.



Special Handlings on Backup Ingress

Relay-Message Method:

1. After receiving **Path with INGRESS_PROTECTION** from primary ingress, it sends **Resv with INGRESS_PROTECTION** to primary ingress, but does **not reserve any resources**
2. When detecting primary ingress failure, it must **keep the state for the Path** message(s) originally received from the primary ingress, update the message(s) and put the message(s) into the bypass LSP tunnel to the next hop(s) of the primary ingress. It **keeps the state for the Resv** message(s) and update the message(s) such as setting Protection-in-use.
3. When the state for the Path message(s) is to be removed by Path Tear from primary ingress or refresh timer expiration with primary ingress up, it **tears down the backup LSP(s) through calling existing functions**

1. Ingress PE5 sends **Path w/ INGRESS_PROTECTION** to backup ingress PE6 after primary LSP is set up

3. PE6 sends PE5 **Resv w/ INGRESS_PROTECTION**

2. PE6 creates backup LSP to protect PE5 through calling existing functions

 Primary LSP
  Path Msg
  Backup LSP
  Resv Msg

Proxy-Ingress Method:

1. After receiving **Path with INGRESS_PROTECTION** from proxy-ingress, it sends **Path with INGRESS_PROTECTION** to primary ingress, but does **not reserve any resources**
2. After receiving **Resv with INGRESS_PROTECTION** from its next hop (i.e., primary ingress), it considers the case that **no normal label is assigned by its next hop** (i.e., primary ingress). It does **not assign a normal label for the interface from its upstream node** (i.e., proxy ingress or primary ingress), **or write a cross connect** from its upstream node to its next hop. It sends a **Resv with INGRESS_PROTECTION** to its upstream node (proxy-ingress), but does **not reserve any resource**.
3. When detecting primary ingress failure, it can not get any Path messages from its previous hop (i.e., the proxy-ingress or the primary ingress), thus it must **keep the state for the Path** message(s) originally received from the primary ingress, update the message(s) and put the message(s) into the bypass LSP tunnel to the next hop(s) of the primary ingress. It can not send any Resv message(s) to its previous hop (i.e., the proxy-ingress or the primary ingress), thus it **keeps the state for the Resv** message(s) originally received from its next hop (i.e., primary ingress) and update the message(s) such as setting Protection-in-use.
4. When the state for the Path message(s) is to be removed by Path Tear from primary ingress or refresh timer expiration with primary ingress up, it **tears down the backup LSP(s) through reusing FRR with changes**

1. Proxy Ingress PE5' (i.e., primary ingress PE5 acting as PE5') sends **Path w/ INGRESS_PROTECTION** to backup ingress PE6

2. Backup ingress PE6 sends **Path w/ INGRESS_PROTECTION** to primary ingress PE5

3. PE5 sends **Path w/o INGRESS_PROTECTION** to NHs (NHs send **Resv** to primary ingress PE5)

6. Backup ingress PE6 sends **Resv w/ INGRESS_PROTECTION** to proxy ingress PE5' (i.e., PE5 acting as PE5')

LSP Path (ERO):
PE5'—PE6—PE5—NH

4. Primary ingress PE5 sends **Resv w/ INGRESS_PROTECTION** to backup ingress PE6 after receiving Resv from NHs


5. PE6 creates backup LSP to protect ingress PE5 **reusing FRR with changes** to existing procedures

 Primary LSP
  Path Msg
  Backup LSP
  Resv Msg

Backup LSP Creation

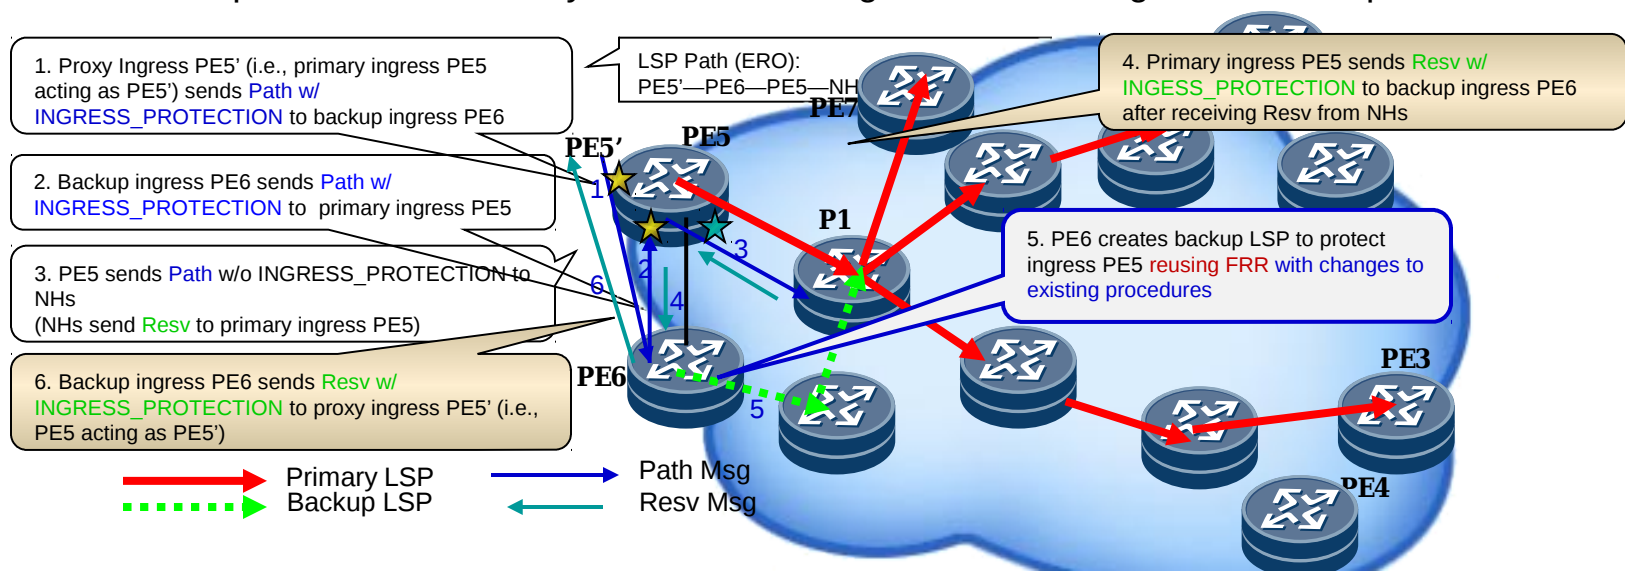
Relay-Message Method:

- Backup LSP creation through calling backup LSP creation function and LFIB entry writing function to import traffic from source into backup LSP.



Proxy-Ingress Method:

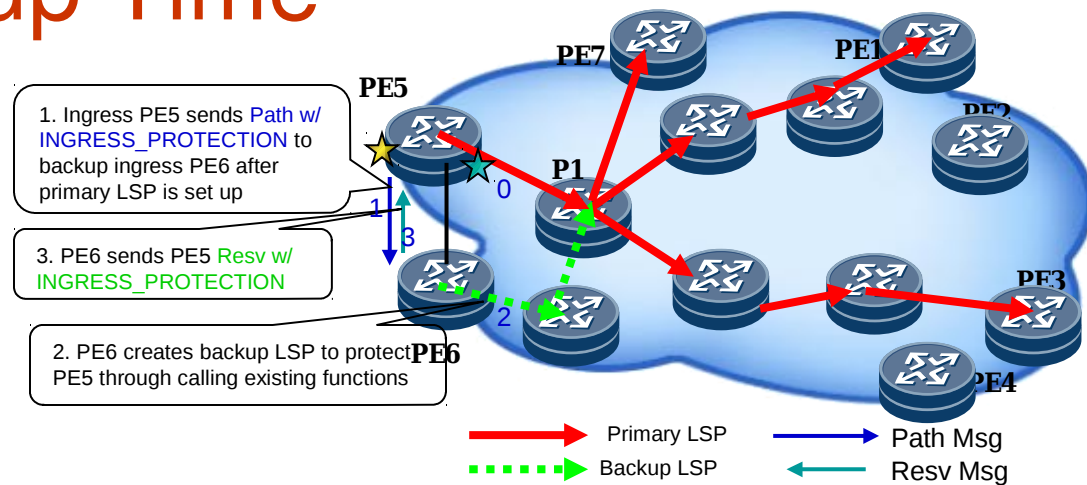
- Backup LSP creation reusing FRR with changes to existing procedures. The changes include: 1) writing LFIB entry to import traffic from source into backup LSP instead of writing LFIB entry to swap incoming label to outgoing label in date packet from upstream hop (i.e., proxy-ingress or primary ingress, when primary ingress fails, traffic from source is imported to backup LSP on backup ingress); 2) getting around the procedure that sets a flag in the LFIB entry indicating inactive. The flag is used to activate the LFIB entry after failure of primary ingress is detected by backup ingress. When source detect is used, the LFIB entry for backup LSP is active always and thus setting the inactive flag needs to be prevented.



Primary LSP Setup Time

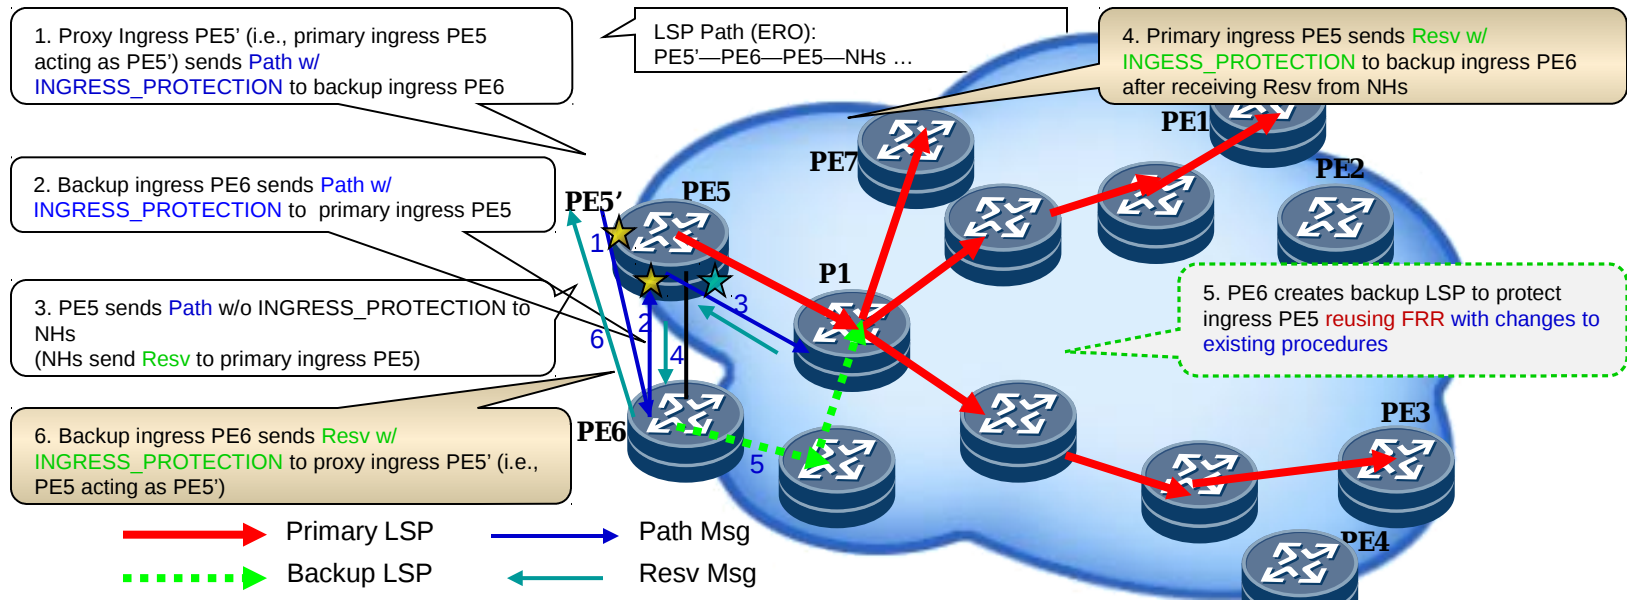
Relay-Message Method:

- Primary LSP setup time after adding ingress protection is the **same as before** (i.e., without adding ingress protection). The setup of primary LSP is not touched after adding ingress protection.



Proxy-Ingress Method:

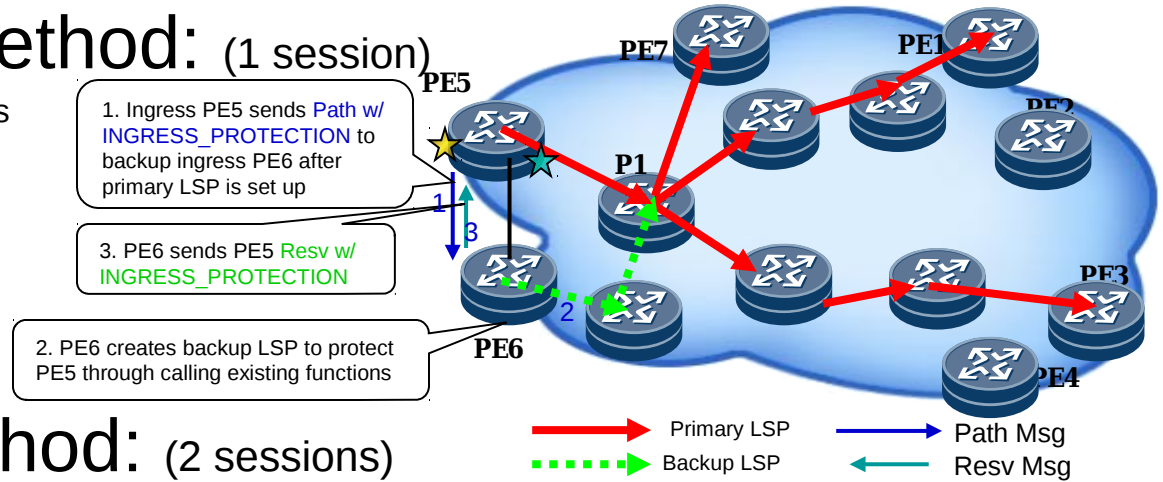
- Primary LSP setup time is **longer** after adding ingress protection. After adding ingress protection, the signalling of primary LSP is changed and goes through two extra hops via backup ingress (i.e., from proxy-ingress (i.e., primary ingress) to backup ingress, from backup ingress to primary ingress, and then from primary ingress to next hops of primary ingress).



Session Maintenance on Primary Ingress

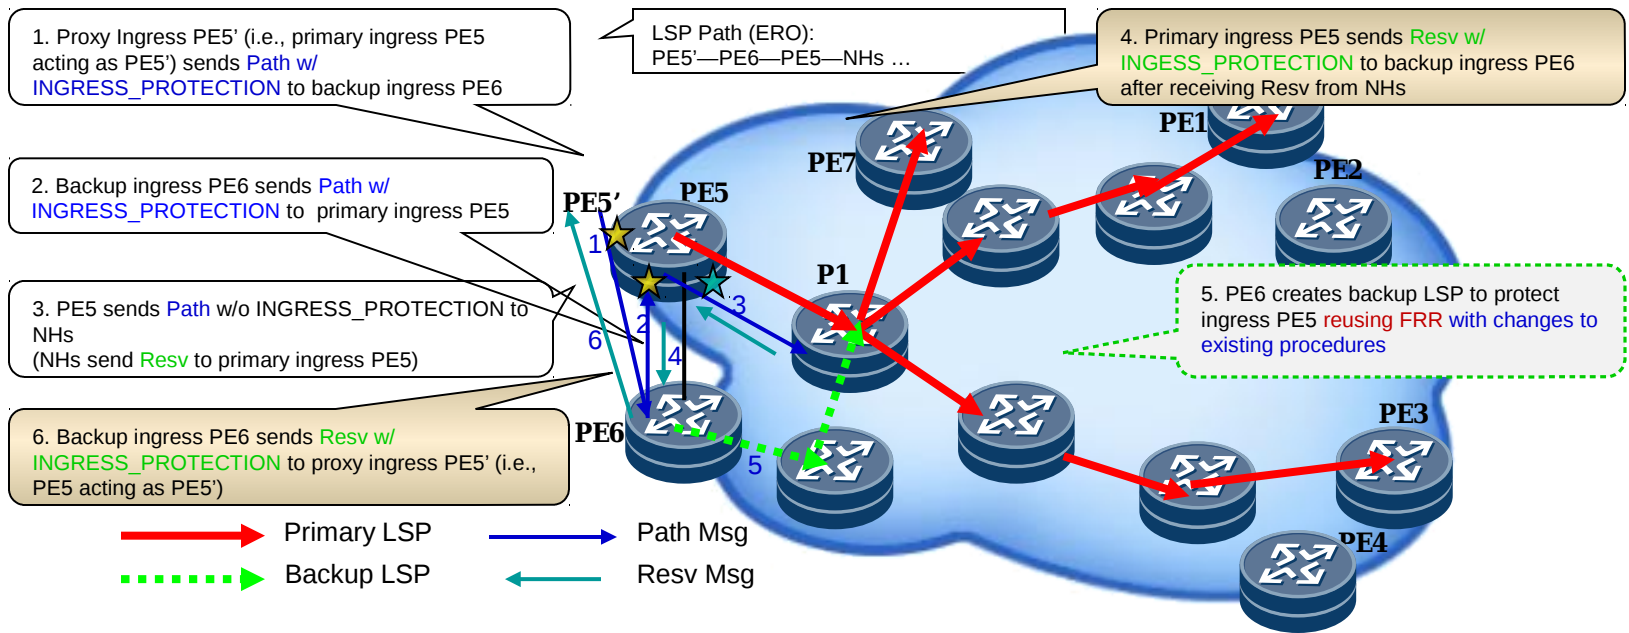
Relay-Message Method: (1 session)

1. Session (state) to backup ingress (refer to 1 and 3)



Proxy-Ingress Method: (2 sessions)

2. Session (state) to backup ingress from proxy-ingress (i.e., "primary ingress") (refer to 1 and 6)
3. Session (state) from backup ingress (refer to 2 and 4)



Session Maintenance on Backup Ingress


Relay-Message Method: (1 session)

1. Session (state) from primary ingress (refer to 1 and 3)

1. Ingress PE5 sends **Path w/ INGRESS_PROTECTION** to backup ingress PE6 after primary LSP is set up

3. PE6 sends PE5 **Resv w/ INGRESS_PROTECTION**

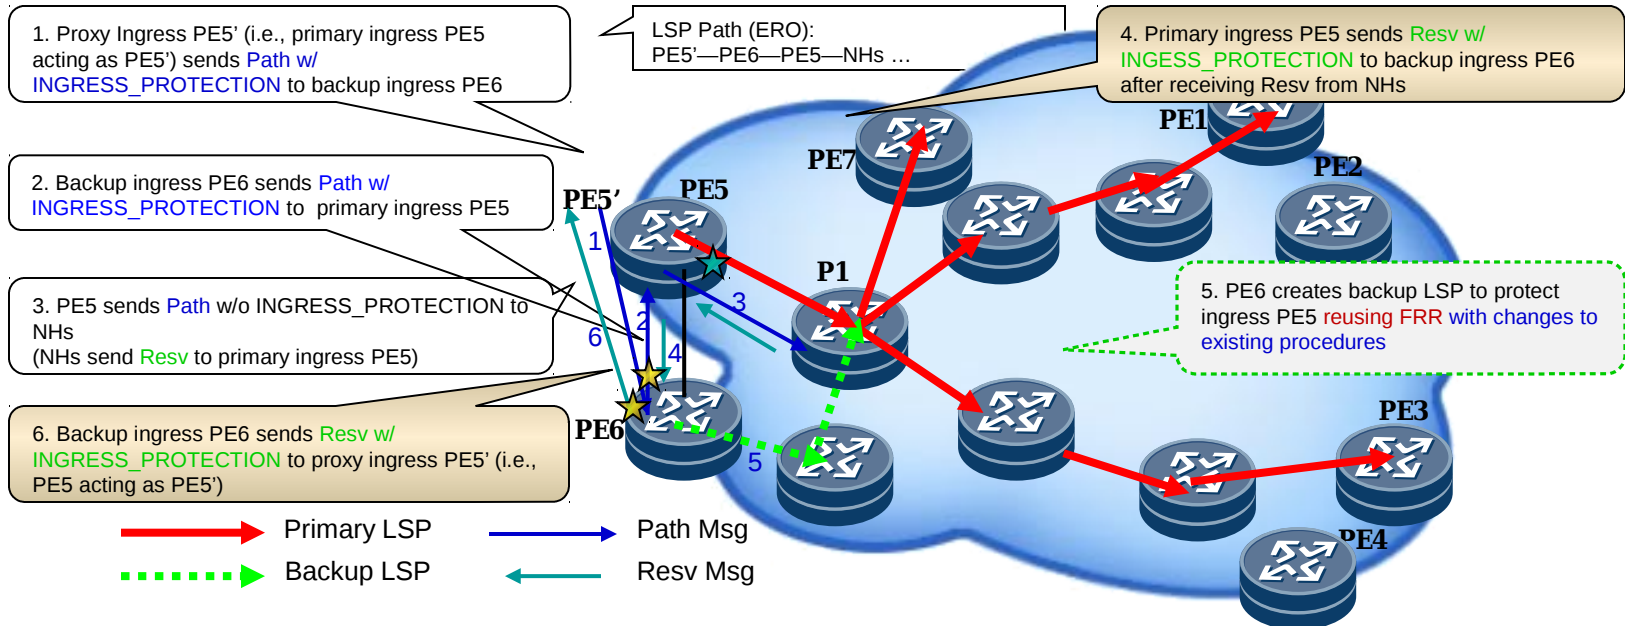
2. PE6 creates backup LSP to protect PE5 through calling existing functions



Proxy-Ingress Method: (2 sessions)

2. Session (state) from (upstream node) proxy-ingress (i.e., "primary ingress") (refer to 1 and 6)
3. Session (state) to (next hop) primary ingress (refer to 2 and 4)

→ Primary LSP
→ Backup LSP
→ Path Msg
← Resv Msg



Scalability

Relay-Message Method: (2 messages, 2 session states)

1. 2 session states
 - one state for session on primary ingress (from primary ingress to backup ingress)
 - one state for session on backup ingress (from primary ingress to backup ingress)
2. 2 messages
 - Path Message with INGRESS_PROTECTION is sent from primary ingress to backup ingress
 - Resv Message with INGRESS_PROTECTION is sent from backup ingress to primary ingress

Proxy-Ingress Method: (4 messages, 4 session states)

3. 4 session states
 - 2 session states on primary ingress (one state for session from proxy-ingress (i.e., “primary ingress”) to backup ingress; one state for session from backup ingress to primary ingress)
 - 2 session states on backup ingress (one state for session from proxy-ingress (i.e., “primary ingress”) to backup ingress; one state for session from backup ingress to primary ingress)
4. 4 messages
 - Path Message with INGRESS_PROTECTION is sent
 - from proxy-ingress (i.e., primary ingress) to backup ingress and
 - from backup ingress to primary ingress
 - Resv Message with INGRESS_PROTECTION is sent

Summary

Item \ Method	Relay-Message	Proxy-Ingress
Configurations	Less	More
Primary LSP depends on backup ingress somehow	No	Yes
Control Message Overhead	Less (2)	More (4)
Special Handlings on Primary Ingress	Less	More
Special Handlings on Backup Ingress	Less	More
Backup LSP creation and deletion	Through calling existing functions	Reusing FRR with changes to existing procedures
Primary LSP setup time after adding ingress protection	Same as before	Longer
Session State Overhead	Less (2)	More (4)
Scalability	Higher(less overheads)	Lower(more overheads)
Local Repair, Global Repair, On-Path Procedure, Revert to Primary Ingress, Backwards Compatibility, Security Considerations	Common for two methods (with minor difference for some cases)	

Thanks

Local Repair on Backup Ingress

After detecting primary ingress failure, backup ingress does the followings

Relay-Message Method:

1. must keep the Path message(s) originally received from the primary ingress, update the message(s) and put the message(s) into the bypass LSP tunnel to the next hop(s) of the primary ingress.
2. keep the Resv message(s) and update the message(s) such as setting Protection-in-use.
3. keep session from primary ingress (inactive)

Proxy-Ingress Method:

4. must keep the Path message(s) originally received from the primary ingress, update the message(s) and put the message(s) into the bypass LSP tunnel to the next hop(s) of the primary ingress.

(since it can not get any Path messages from its previous hop (i.e., the proxy-ingress or the primary ingress),

6. keep the Resv message(s) originally received from its next hop (i.e., primary ingress) and update the message(s) such as setting Protection-in-use.

(since it can not send any Resv message(s) to its previous hop (i.e., the proxy-ingress or the primary ingress)

8. keep session from (upstream node) proxy-ingress (i.e., “primary ingress”) (inactive)
9. Let session to (next hop) primary ingress down

1, 2 and 3 are common for two methods (note 3 considered to be common)

In summary, common for two methods with minor difference

Global Repair

After detecting primary ingress failure, backup ingress does the followings for global repair
(common for two methods)

Relay-Message Method:

Proxy-Ingress Method:

- Compute a global optimal path
- Set up a new LSP along the path with a different LSP ID
- Switch traffic to the new LSP
- Tear down the old LSP

Backwards Compatibility

Primary ingress sends Path message with INGRESS_PROTECTION to backup ingress.
If backup ingress does not support the extensions for ingress protection, then

Relay-Message Method:

- Primary ingress receives Path Err from backup ingress and sends Path Tear to backup ingress

Proxy-Ingress Method:

- Proxy-ingress (i.e., “Primary ingress”) receives Path Err from backup ingress and sends Path Tear to backup ingress
- Primary ingress changes the ERO and sends Path message with the updated ERO to its next hops after determining that backup ingress does not support ingress protection

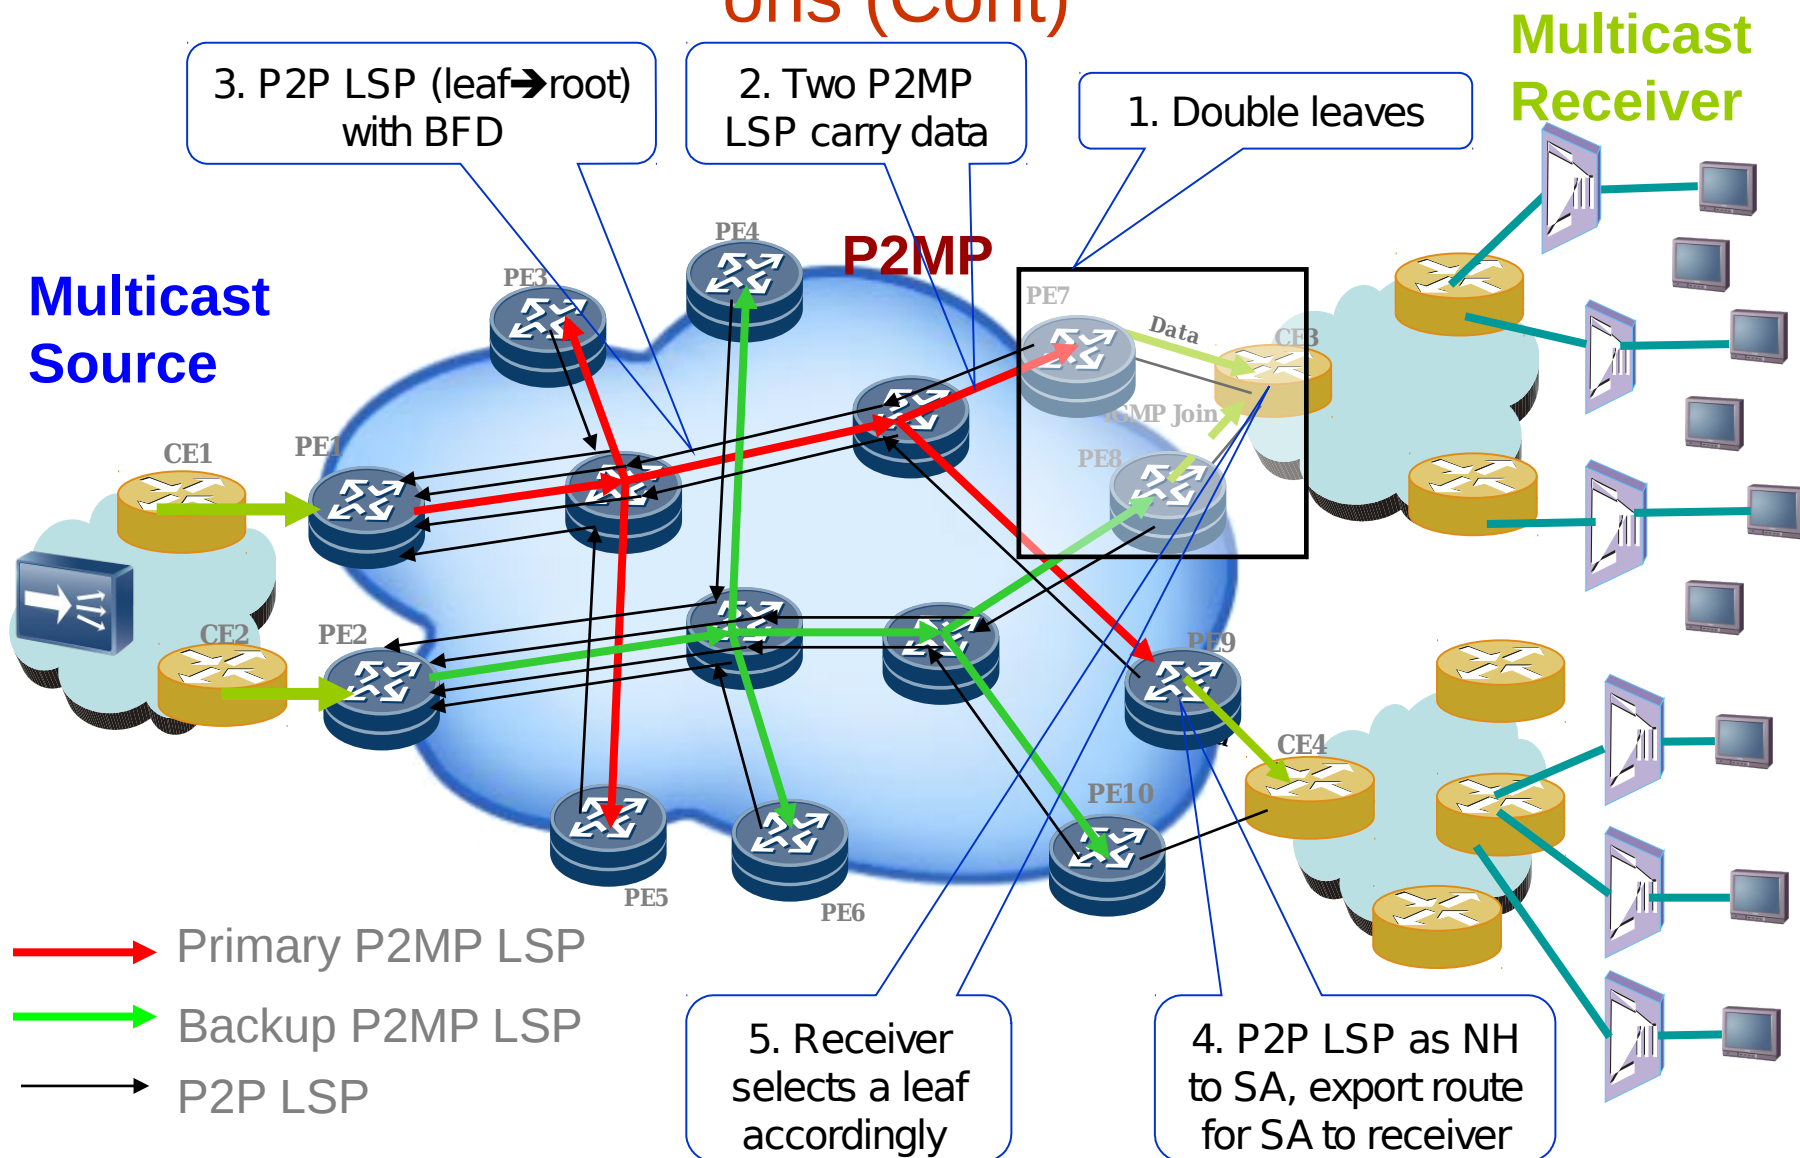
common for two methods with minor difference

Details in Existing P2MP LSP Ingress & Egress Protections

➤ To provide E2E P2MP LSP protection , a **current way (detail in next page)**

- Redundant Root and Every Leaf
- Create two P2MP LSPs from root to leaves, carry the same data at same time.
- For each leaf, create a P2P LSP from the leaf to root and configure BFD with it
- Run iBGP on every leaf node and use P2P LSP as its next hop
- When BFD detects P2P LSP failure, BGP withdraws route to root and this makes the receiver switch to another leaf to get the data.

Details in Existing P2MP LSP Ingress & Egress Protections (Cont)



Issues in Existing P2MP LSP Ingress & Egress Protections

➤ Not easy to operate

- For each P2MP LSP branch/sub-LSP,
 - need configure a reverse P2P LSP from leaf to root with BFD
 - P2P LSP with BFD is used to detect failure of its corresponding P2MP sub-LSP

➤ Not reliable

- The failure of reverse P2P LSP from leaf to root does not mean the failure of its corresponding P2MP sub-LSP from root to leaf

➤ Consume lots of resource

- Reserve/use double bandwidth

➤ Speed of Global Recovery

- Depends on convergence of IGP and BGP

Advantages of P2MP LSP Ingress and Egress Local Protection

- All parts of P2MP LSP are locally protected
- Only one P2MP LSP is used to implement an E2E protection
 - Normally two P2MP LSPs are used
- Big saving on resource : 50% bandwidth saving
 - No need to reserve/use double bandwidth
- Faster recovery
 - Speed of local protection recovery
 - Flow recovery within 50ms when a failure happens
- Easier to operate