



Flow Recovery Support for ECS Rules

This chapter describes the Flow Recovery feature and provides detailed information on the following topics:

- [Feature Description, on page 1](#)
- [How It Works, on page 5](#)
- [Configuring Flow Recovery Checkpointing, on page 6](#)
- [Monitoring and Troubleshooting the Flow Recovery Feature, on page 8](#)

Feature Description

The Flow Recovery feature is introduced to support flow checkpoint for ECS rules. Flows of a rule can be recovered if the rule is made eligible. To make any rule eligible for checkpointing or recovery, the rule must be configured in the service scheme framework and identified based on rulename. Any type of rule can be eligible for checkpointing irrespective of the rule definition or protocol type. The rule type can be static/predefined, group-of-ruledef, dynamic, or ADC rule.



Important Flow Recovery is a licensed Cisco feature requiring a separate feature license. Contact your Cisco account representative for more information.

With the previous implementation, Cisco PCEF supports ICSR/SR check pointing at session level. Session and bearers were recovered as part of session recovery or ICSR switchover but flows were not recovered. Due to this when the packets were received on already established flows after the recovery, the packets were matched to default rules based on L3/L4 definition. This causes billing impact as flows were incorrectly charged.



Important If the flow recovery/checkpointing feature is enabled, then ICSR control outage might have a significant impact.

The enhancements with the Flow Recovery feature are listed below.

- **List of eligible rules for flow recovery:** A configurable trigger action CLI **flow recovery** is added to configure flow-recovery under service scheme. The flows matching the configured rules in service scheme will be recovered. The flows will be checkpointed to AAAMgr when the delay timer for the flow matching these rules expires post the flow establishment.

- **Limit configuration for the number of flows recovered:** The maximum number of flows to be checkpointed or recovered will be limited per subscriber using the **system-limit flow-chkpt-per-call** CLI at active-charging service level. This limit controls the number of flows that can be recovered across rules and bearers for a single call.

The number of flows checkpointed/recovered will be limited at per SessMgr instance level. This will be a non-configurable value.

- **Rule matching post recovery:** If a flow is recovered to match a checkpointed rule, the incoming packets will continue matching the same rule throughout the lifetime of the flow and the flow will remain checkpointed with the rule. The checkpoint information will be deleted post recovery when:
 - The rule from recovered list is uninstalled for the session
 - The flow idle time-out happens
 - The TCP-based application flow is gracefully terminated when a RST packet is received at the gateway
 - The TCP-based application flow is not terminated when FIN packet is received at the gateway and cleared after the flow idle timeout happens

- **Handoff/ICSR scenario:** After a rule is recovered for control events such that the bearer/rule information is not available (due to delay in EGTP message), the packets for these flows will match the default rule based on L3/L4 definition until transactions are complete. Once the control events are no longer pending, the packets will start matching the last matched rule from flow recovery list.

Similarly for ICSR, the packets for recovered flows will match the default rules based on L3/L4 definition. Once the information is received from new chassis, the packets will start matching the last matched rule from flow recovery list.

- **KPI and Bulk Statistics:** KPI and Bulk Statistics are added on a per rulename basis for the rules that are eligible for checkpointing. When a rule is configured as eligible for recovery, KPI and bulk statistics for the rule will be displayed. For a rule that is deleted from the service-scheme framework, KPI and bulk statistics will not be displayed. SRP level statistics are also added to maintain the SRP bandwidth required for flow checkpoint.

Relationships to Other Features

This sections describes the relationship of the Flow Recovery feature with other supported features.

Feature Interaction

The following table describes the impact to inline services with the flow recovery implementation.

The behavior applies to recovered flows unless explicitly mentioned.

Inline Services	Flow Recovery Impact	Description
CF Static and Dynamic	No	CF redirect and CF blacklist post recovery will not work and packets will be allowed.
ICAP	No	Post recovery GET messages for the recovered flow will not be forwarded to ICAP server.

Inline Services	Flow Recovery Impact	Description
Tethering Detection	No	For IP-TTL based tethering detection, the post recovery flow will be considered as tethered if the first packet post recovery is uplink packet. For OS-DB signature based tethering detection, the post recovery flow will be considered as tethered if the SYN packet received post recovery has matching signature in the database.
P2P	No	P2P rules are not supported for recovery.
Post processing rules	No	URL Redirect for recovered flows will not be supported based on post processing rules. The default action "Discard" will be applied.
X-Header insertion	No	X-header insertion will not work for recovered rules.
DSCP/IP-TOS Marking	Yes	Packets will be marked post recovery with correct DSCP/IP-TOS values.
Idle-timeout handling	Yes	The flow is terminated after idle timeout and checkpoint limits are decremented.
TCP based application Protocol	No	Analyzer will not be engaged for recovered flows.
UDP Based application Protocol	Yes	Analyzer will be engaged for recovered flows if configured in rulebase.
ICSR	Yes	Flows will be recovered to match the last match rule pre ICSR and for which checkpoint was successful.
Session Recovery	Yes	Flows will be recovered to match the last match rule pre SR and for which checkpoint was successful.
IPv6	Yes	IPv6 flows will be recovered.

Flow Actions

The following table describes the behavior of flow actions with the flow recovery implementation.

This applies to recovered flows unless explicitly mentioned.

Flow Action	Description
Rate Limiting	Per bearer rate limiting, per rule rate limiting, ITC rate limiting, or QG related rate limiting is applied.

Flow Action	Description
Allow/Deny packet	Flow-action discard OR dynamic-rule flow-status enforcement is applied.
URL redirect	Run time URL redirection is not applied for a recovered flow. If the packets of a flow are redirected before recovery and the matching rule is eligible and checkpointed, packets are redirected post recovery as well.
Flow Readdress	Since the flow readdress is identified at SYN packet and SYN packets are not analyzed post recovery, flow readdress will not work post recovery.
Flow Terminate	Flow terminate action is applied for dynamic changes to charging action configuration.
Next Hop	Packets are forwarded to the next hop address.

Charging Methods

The following table describes the impact to ECS charging methods with the flow recovery implementation.

The behavior applies to recovered flows unless explicitly mentioned.

Charging Method	Flow Recovery Impact
eGCDR	No
PGW-CDR	No
VoGx	No
Gy	No
RF	No
EDR	In case of recovery, generated EDR will show "Unknown" for sn-direction because the direction of first packet is not known.
FDR	No
UDR	No

Limitations

This section lists the limitations associated with the Flow Recovery/Checkpoint feature.

- Few calls may be lost if session recovery occurs when the system is experiencing very high flow setup rates.
- Checkpoint failures (failure counters being incremented) might be seen when large number of flows are being check pointed (>12K per seconds). The failed checkpoints are retried again. If a session recovery, unplanned card migration, or unplanned switchover occurs while the checkpoints are being retried, call loss may be seen.
- Disabling of flow recovery feature (via CLI) while active calls (using flow recovery) are present might result in session manager restart.

- Rule Matching – Packets are not matched to L7 rule after bearer movement from dedicated bearer to default bearer during LTE to WiFi HO. The same behavior will be seen post recovery.
- Static rule match on dedicated bearer – After recovery, flow packets will always match the last matched rule. If the TFT are installed on dedicated bearer post recovery, such that the last match checkpointed rule is a static rule but the packets are received on dedicated bearer due to matching TFT, the packets are still matched to static rule but on dedicated bearer. Charging and QoS parameters are applied as per dedicated bearer.
- EDR loss for SR/Unplanned ICSR – EDR is not generated in case of session recovery or unplanned ICSR.
- Parent-child relation loss post recovery – The parent-child relationship post recovery is not maintained for protocols such as RTP and RTSP.
- Lifetime of recovered flows – Lifetime will not be incremented for recovered flows for which no packet is received post recovery and is deleted due to flow idle timeout.
- Service scheme framework – Flow checkpointing is dependent on the service scheme framework. Since this framework is introduced only in release 20.2, flow checkpointing for the existing calls/flows is supported only for the 20.2 build during upgrade scenarios. In all other upgrade scenarios, flow checkpointing will not be supported for the existing calls/flows.

How It Works

This section describes the Flow Recovery configuration. The Service Scheme framework configuration is required to configure and enable this feature for a subscriber.

Use the sample configuration to enable the Flow Recovery feature.

```
configure
  active-charging service s1
    trigger-action ta1
      flow-recovery
    #exit
    trigger-condition tc1
      rule-name = rule01
      delay = 600
    #exit
    service-scheme ss1
      trigger flow-create
        priority 1 trigger-condition tc1 trigger-action ta1
      #exit
    #exit
end
```

Restrictions

Any TCP-based application protocol sends out mid-flow acknowledgements that are TCP packets. These packets are generally matched to a L3/L4 based rule on that bearer. In such cases a delete checkpoint will be sent whenever the L3/L4 rule is processed for the flow and again a checkpoint to AAAMgr is sent when the L7 rule is processed for the flow. This could result in high number of checkpoints sent/received and delete

checkpoints sent/received. To avoid this scenario, delay charging must be configured for all packets or mid-flow packets in conjunction with flow-recovery for the rule under rulebase using the **flow control-handshaking charge-to-application { all-packets | mid-session-packets }** CLI configuration.

Configuring Flow Recovery Checkpointing

The rules eligible for checkpointing must be configured in the service-scheme framework and identified based on rulename. The rule type can be static/predefined, group-of-ruledef, dynamic or ADC rule. A group-of-ruledef can be independently made eligible for checkpointing even if the constituent ruledefs are not eligible for checkpoint. If both group-of-ruledef and constituent ruledefs are eligible for checkpoint, either one can be checkpointed based on the higher action priority configured in rulebase.

A delay timer to delay the checkpoint of a flow can be configured under the service-scheme trigger condition. This configuration controls the checkpoint sent/received (both checkpointing and deleting checkpoint) for flows which may last for a shorter duration. A new trigger action type "flow recovery" is also added to the service-scheme configuration.

Configuring Flow Recovery

Use the following configuration in the ACS Trigger Action Configuration mode to enable flow recovery for a trigger-action.

```
configure
  active-charging service <service_name>
    trigger-action <trigger_action_name>
      [ no ] flow-recovery
    exit
```

Notes:

- The flows for the rule will be checkpointed as per session level and call level limit.
- To disable flow recovery, configure the following command:

```
no flow-recovery
```

Configuring Delay Checkpointing for flow

Use the following configuration in the ACS Trigger Condition Configuration mode to define the list of rules along with the delay after which the flows can be checkpointed. When configured in conjunction with flow-recovery trigger, the flows for the rule(s) will be checkpointed as per session level and call level limit after the delay timer is expired.

```
configure
  active-charging service <service_name>
    trigger-condition <trigger_condn_name>
      delay = <delay_time>
      rule-name operator <rule_name>
    exit
```

Notes:

- The range of the delay timer to be configured is 1 through 600 seconds. If the "delay" CLI command is not configured under trigger-condition, any flow for the rule will be checkpointed immediately on flow creation.
- **rule-name operator** <rule_name>
 - *operator* specifies to match and must be one of the following:
 - =
 - *contains*
 - *ends-with*
 - *starts-with*

The operators "contains", "starts-with", and "ends-with" cannot be used with dynamic rule names. For dynamic rules, the entire rulename must be mentioned with the "=" operator.

 - *rule_name* must be an alphanumeric string of 1 through 63 characters.
- In any defined trigger-condition, a user can configure upto a maximum of 15 entries.
- To have more rules eligible for flow checkpoint, a user can configure multiple trigger condition(s) associated with the same trigger-action.
- Use the **no rule-name** command to remove the particular rule from the list of eligible rules for flow checkpoint. For wildcard-based rule definition, this command must contain rulename in the same format.
- Use the **no delay** command to checkpoint all eligible rules immediately without any delay.

Enabling trigger for new flows

Use the following configuration in the ACS Service Scheme Configuration mode to enable trigger for every new flow.

```
configure
  active-charging service <service_name>
    service-scheme <serv_scheme_name>
      [ no ] trigger flow-create
    exit
```

Configuring Flow Checkpoint Limit

Use the following configuration in the ACS Configuration mode to control the number of flows that can be checkpointed per call.

```
configure
  active-charging service <service_name>
    system-limit flow-chkpt-per-call <num_flows>
  exit
```

Notes:

- The number of flows that will be allowed to checkpoint will be maintained at call-ID level. The number of flows that can be checkpointed per call are between 1 to 300.
- The default value of the number of flows checkpointed is 10. This value will apply for all the calls across rules and bearers.

Configuring Flow KPI Bulk Statistics

A new bulk statistic schema **flow-kpi** is added to maintain the KPI for flows. These counters will be the sum of the statistics for the corresponding rule(s) across session manager instance.



Important The Flow Recovery feature license is required to display KPI and configure bulk statistics.

Use the following configuration to configure the Flow KPI bulk statistic schema.

```
configure
  bulkstat mode
    file file_number
    flow-kpi schema schema_name format schema_format
  exit
```

Use the following configuration to view the Flow KPI bulk statistics in the Flow KPI schema.

```
show bulkstats variables flow-kpi
```

Verifying the Flow Recovery Configuration

Enter the following command to check the cumulative KPI for each rule across session managers:

```
show active-charging flow-kpi [ all | instance instance_id ]
```

Monitoring and Troubleshooting the Flow Recovery Feature

The sections listed below describe the KPI, Bulk statistics, and SRP level checkpointing statistics that can be used to identify if the checkpoints are sent and successful.

Flow Recovery Show Command(s) and/or Outputs

This section provides information regarding show commands and/or their outputs in support of the Flow Recovery feature.

show active-charging flows all

The output of this command is modified to include a new flag to identify if the flow is a recovered flow or not. The value will display Y for a recovered flow and N for a non-recovered flow. The flag will be displayed if flow recovery license is present.

- Recovered Flow: (Y) - Yes (N) - No

show active-charging flow-kpi all

In support of the Flow Recovery feature, the **show active-charging flow-kpi all** CLI command is added to display the cumulative KPI for each rule across session managers and can also filter the statistics for all rules based on sessmgr instance. KPI are added on a per rule basis.

- Rule Name
- Active Flows
- SR Flow Checkpoint Sent
- SR Flow Checkpoint Received
- GR Flow Checkpoint Sent
- GR Flow Checkpoint Received
- SR Flow Checkpoint Delete Sent
- SR Flow Checkpoint Delete Received
- GR Flow Checkpoint Delete Sent
- GR Flow Checkpoint Delete Received
- Flows of lifetime bucket1
- Flows of lifetime bucket2
- Flows of lifetime bucket3

show active-charging trigger-action all

The following field indicates whether flow recovery is enabled or disabled.

- Flow recovery

show active-charging trigger-condition all

The following fields are added to the output of this command in support of this feature:

- Trigger Action Delay - Displays the delay (in seconds) for application of action.
- Rule-name/GOR - Displays the condition specified for a particular rule/GoR for flow checkpoint.

show srp checkpoint statistics

The SRP checkpoint statistics are added as part of checkpoint and recovery for ECS and ADC flows. Statistics for any flow checkpoint sent and received as well as delete checkpoint sent and received will be maintained.

ACS_FLOW_INFO and DEL_ACS_FLOW_INFO fields are displayed only after the creation/deletion of flow checkpoints respectively. When the flow checkpoints are sent from active chassis to standby chassis, the corresponding ACS_FLOW_INFO micro checkpoint statistics will be incremented. Similarly, after the expiry of idle timeout configured in the active-charging service, the flows will be deleted and the corresponding DEL_ACS_FLOW_INFO micro checkpoint statistics will be incremented.

A sample output is shown below. The rate limited MC checkpoints are displayed under "Checkpoints Rate limited" statistics and will be pegged only for ACS_FLOW_INFO and DEL_ACS_FLOW_INFO micro checkpoints.

```

show srp checkpoint statistics [OR] show srp checkpoint statistics sessmgr all
SRP Session Checkpoint Sent Statistics
=====
Checkpoints Rate limited-----+
Bytes sent-----+
External_audit_in_progress-----+
Recovery in progress-----+
Write list congested-----+
Invalid GR state-----+
Session not ready for recovery-----+
Encoding failure ignored-----+
Encoding failure-----+
Ignored as Invalidate CRR sent-----+
Ignored as first FC never sent-----+
Checkpoint rate (last 30 min)-----+
Checkpoint rate (now)-----+
Encoding attempted-----+
Messages sent-----+
Checkpoint-type      | | | | | | | | | | | | | |
=====
UPDATE_CLPSTATS      0  0  0  0  0  0  0  0  0  0  0  0  0  0  3  0  0
ACS_SESS_INFO        2060 2060 0  0  0  0  0  0  0  0  0  0  0  0  298130 0
ACS_CALL_INFO        2060 2060 0  0  0  0  0  0  0  0  0  0  0  0  306240 0
PGW_UBR_MBR_INFO    1000 1000 0  0  0  0  0  0  0  0  0  0  0  0  457000 0
ACS_FLOW_INFO        30  50  0  0  0  0  0  0  0  0  0  0  0  0  8490 20
DEL_ACS_FLOW_INFO   30  50  0  0  0  0  0  0  0  0  0  0  0  0  6060 20
Full Checkpoint     1000 1000 0  0  0  0  0  0  0  0  0  0  0  0  4754470 0
=====

```

Flow Recovery Bulk Statistics

A new **Flow KPI** schema is added in support of the Flow Recovery feature.

The following bulk statistics are added in the Flow KPI Schema:

- ecs-flow-rule-name – The name of rules eligible for flow checkpointing.
- ecs-num-active-flow – Total number of active flows of the rule.
- ecs-flow-chkpt-sent-sr – Total number of SR flow checkpoint sent for the rule.
- ecs-flow-chkpt-recvd-sr – Total number of SR flow checkpoint received for the rule.
- ecs-flow-chkpt-sent-gr – Total number of GR flow checkpoint sent for the rule.
- ecs-flow-chkpt-recvd-gr – Total number of GR flow checkpoint received for the rule.
- ecs-del-flow-chkpt-sent-sr – Total number of SR delete flow checkpoint sent for the rule.
- ecs-del-flow-chkpt-recvd-sr – Total number of SR delete flow checkpoint received for the rule.
- ecs-del-flow-chkpt-sent-gr – Total number of GR delete flow checkpoint sent for the rule.
- ecs-del-flow-chkpt-recvd-gr – Total number of GR delete flow checkpoint received for the rule.
- ecs-flow-lifetime-bucket1 – Total number of flows of lifetime_bucket1. Lifetime value of bucket1 is configurable.

- ecs-flow-lifetime-bucket2 – Total number of flows of lifetime_bucket2. Lifetime value of bucket2 is configurable.
- ecs-flow-lifetime-bucket3 – Total number of flows of lifetime_bucket3. Lifetime value of bucket3 is configurable.

