November 2021

MECHANISM TO QUERY USAGE FROM USER PLANE IN OPTIMISED AND RELIABLE MANNER

Mukesh Yadav

Ravi Shankar Mantha

Sandeep Dasgupta

Milind Nadkarni

Sanjeev Panem Jaya

MECHANISM TO QUERY USAGE FROM USER PLANE IN OPTIMISED AND RELIABLE MANNER

AUTHORS:
Mukesh Yadav
Ravi Shankar Mantha
Sandeep Dasgupta
Milind Nadkarni
Sanjeev Panem Jaya

ABSTRACT

Within a 3rd Generation Partnership Project (3GPP) fifth generation (5G) telecommunications environment, various Usage Reporting exchanges may entail different inefficiencies. To address such inefficiencies, techniques are presented herein that optimize Usage Reporting exchanges. In particular, aspects of the presented techniques support the introduction of a new custom information element (IE) on the Packet Forwarding Control Protocol (PFCP) interface. The custom IE enables a user plane (UP) to selectively report multiple Usage Reporting Rules (URRs) that satisfy a category of URR as specified in the IE. Aspects of the presented techniques provide a mechanism involving an optimized PFCP modification request packet that obviates the need to send hundreds of URR identifiers in a request. Additionally, aspects of the presented techniques may also be used within a 3GPP fourth generation (4G) telecommunications environment.

DETAILED DESCRIPTION

Within a 3rd Generation Partnership Project (3GPP) fifth generation (5G) telecommunications environment, charging is achieved in collaboration with network functions such as a Policy Charging Function (PCF), a Session Management Function (SMF), a User Plane Function (UPF), and a charging function (CHF).

Briefly, an SMF configures a UPF with usage rules – i.e., Usage Reporting Rules (URR) on an N4 interface utilizing the Packet Forwarding Control Protocol (PFCP) protocol. The UPF performs usage reporting to the SMF through PFCP message exchanges. The SMF relays the received information to a CHF in the Used Unit Container element of a Charging Data Request message.

An exchange between a UPF and an SMF may be completed using either asynchronous reporting or synchronous reporting.

For asynchronous reporting, usage reporting occurs by the UPF and is relayed to the SMF in a PCFP session report message exchange. When the SMF receives the usage reports it sends them to the CHF. The UPF sends the usage report whenever any of the triggers that are configured by the SMF are met.

For synchronous reporting, change condition (CC) events (such as, for example, AMBR_CHANGE, PLMN_CHANGE, etc.) are provisioned by the CHF and are detected at the SMF. In this case, the SMF queries the usage reports of all of the URR that are associated with an event by sending a PFCP Modification Request message to the UPF and receiving a usage report from the UPF in response. Such events are called change in charging condition events.

Aspects of the techniques presented herein provide enhancements for synchronous reporting. In particular, when a CC event arises, the SMF needs to query the URRs for their current usage. The relevant standard (i.e., section 7.5.4.1 of the 3GPP Technical Specification (TS) 29.244) provides for two mechanisms to accomplish such an inquiry.

A first mechanism encompasses a query list containing the URR identifiers (IDs) for which usage is required. With this approach, the user plane (UP) will relay the usage of all of the URRs that are queried. A second mechanism encompasses sending a Query All URRs (QAURR) flag in a PFCP Session Modification Request message. With this approach, a UP element will relay the usage for all of the URRs that are created on the UP element.

The second mechanism represents the standard technique for querying all of the URRs when a CP has determined that a CC event encompasses all of the created or active URRs. Such an approach works fine if there is interest in receiving usage particulars for all of the created URRs on a UP element.

In a 5G telecommunications environment, a URR that is created on a control plane (CP) element or on a UP element belongs to a particular category. URR categories may include:

1. A rating group-driven URR, employing the concept of flow-based charging. See, for example, section 5.2.1.5 of the 3GPP TS 32.255.

2. A quality of service (QoS) flow-based URR, employing the concept of QoS flow-based charging. See, for example, section 5.2.1.6 of 3GPP TS 32.255.
3. A session level URR that is created on behalf of session limits.
4. A URR that is created for PCF accounting that employs application URRs.
5. A custom URR that is created for Remote Authentication Dial-In User Service (RADIUS) accounting and that is used within an enterprise solution.

A CC event may be armed by the CHF at the SMF for one or more categories. Such categories may include:

1. Any kind of URR, with arming by session triggers.
2. Flow Based Charging (FBC) -relayed URRs, with arming by MultiUnitInformationRG driven triggers.
3. A QoS flow-based charging (QBC) -related URR.
4. A custom category of URRs on a CP such as, for example, Radius accounting for enterprise deployments.
5. Besides a CC-driven query, a QAURR will query everything. There may be cases where URRs that are created for PCF accounting will not need to be reported or queried when a CC event occurs.

The standards-driven second mechanism that was described above has a limitation if a CC event happens for one of the category and Query usage is not to be done for all URR's for a session. In particular, if querying is performed using a list of URR ID values when an event is armed for all of the URRs, such an approach is not optimized and it will result in a larger packet size. Additionally, the CP can only query the URRs that are known to the CP.

There are three kinds of rules that are possible in a deployment. A first kind of rule encompasses dynamic rules that are fully created by a PCF. Such rules are specification-driven. A second kind of rule encompasses predefined rules that are configured on a CP or a UP and which are activated by a PCF. Such rules are specification-driven.

A third kind of rule encompasses static rules. Such rules are not specification-driven. However, rules of this kind are employed in commercial deployments because of the value that they add. In such an approach, many rules are configured under a common umbrella concept like rules grouping where a grouping name is sent by a CP to a UP during

PFCP session establishment. A benefit of this type of approach is that there are rules which are applicable for a particular session or Data Network Name (DNN). They can be assumed to be present on both the CP and the UP so no activation will be required by a PCF or a SMF, thus obviating a potentially considerable amount of signaling on N7, N4, N1, or N2 interfaces. A problem with the static rule concept is in current scope of Querying usage is stated below.

URRs for a static rule may be created on the SMF in two different ways. A first way encompasses creation when a call is received. A URR context may be created on a CP for all of the possible static rules. However, such an approach is not optimized. For example, 100 or 1,000 such rules may be present and a piece of user equipment (UE) might employ just a few of the rules during the lifetime of a session. Creating a URR for all such rules will unnecessarily consume memory on the CP. Even if the CP does perform all of the creations, a subsequent query of all of the IDs will result in congestion on an N4 interface.

A second way for creating a static rule encompasses creation when a first usage report arrives. At that time, a URR context may then be created. Such an approach raises another problem in that the UP might activate a static rule and begin accounting. However, its first usage may not yet be reported and, thus, the CP may not be aware of its existence.

Considering such issues, the standards-driven second mechanism involving the query through a QAURR flag provides the best approach for querying URRs without sending a list of all of the URRs. However, such an approach also has a limitation.

Like the SMF, the CP may have different categories of URRs (such as the URR categories that were described above). For a CC event that happens at the CP, the event could just be armed for QBC or FBC. Hence, a query must take place accordingly. PCF-driven URRs, session level URRs, and custom URRs do not need to be queried or reported upon. In such cases, the usage of a QAURR will lead to the reporting of all possible categories of URRs when actual reporting is needed just for a selected set of URR belonging to said category.

To address the different problems that were identified in the preceding narrative, techniques are presented herein that optimize various of the query usage exchanges that were described above.

Aspects of the techniques presented herein support the introduction of a new custom information element (IE), referred to herein as the QUERY_INTERFACE IE. Such an IE can be sent along with the QAURR flags that may be set in a N4 Modification Request message if reporting on all of the URRs is not expected.

During operation, the QUERY_INTERFACE IE will only be sent with a QAURR flag if reporting on all of the URRs is not needed. Additionally, the QUERY_INTERFACE IE will not be sent when a query list is provided.

The QUERY_INTERFACE artifact is a composite IE which has allocated bits for various interfaces, similar to a reporting trigger IE (see, for example, Section 8.2.19 of 3GPP TS 29.244). The QUERY_INTERFACE IE bit allocation may comprise:

Octet 1-2: Type indication;

Octet 3-4: Length indication;

Octet 5: Spare, Spare, Spare, Intfc5, Intfc4, Intfc2, Intfc2, Intfc1

With such a structure, one or more bits may be set by the CP. The Octet 5 bits may used as follows:

Intfc1: Maps to offline rating group (RG) -driven URRs.

Intfc2: Maps to online RG-driven URRs.

Intfc3: Maps to QBC-driven URRs.

Intfc4: Maps to a session limits URR.

Spare: May be extended in the future for any other interface.

URRs that are created by the CP or that are created at the UPF itself for static rules or predefined rules have to create URR ID values and associate a proper interface category either within a URR ID or a URR context at the UP. Note, that a URR ID may be created by a CP or by a UP (see, for example, Section 8.2.54 of 3GPP TS 29.244). There are bits reserved for various functionality. Here, an interface ID may be reserved as four of the least significant bits (LSBs) for the CP or the UP when assigning a URR ID for a URR. It must be properly assigned an interface ID. When the UP receives a QAURR along with the QUERY_INTERFACE IE, the UP is to report all of the URRs which satisfy the interface criteria that are identified in the request. For example, if the CP sends a PFCP modification request and a QAURR flag is set to TRUE, then the QUERY_INTERFACE

IE can be set with Intfc1 or Intfc2. Subsequently, the UP will search all of the URRs and will report the URRs whose interface ID corresponds to 1 and 2.

Aspects of the techniques presented herein may be further explained with reference to two exemplary flow diagrams that will be illustrated and described in the narrative that is presented below. A first flow diagram, which is presented in Figure 1, below, depicts elements of a query usage example that employs standards-based mechanisms.

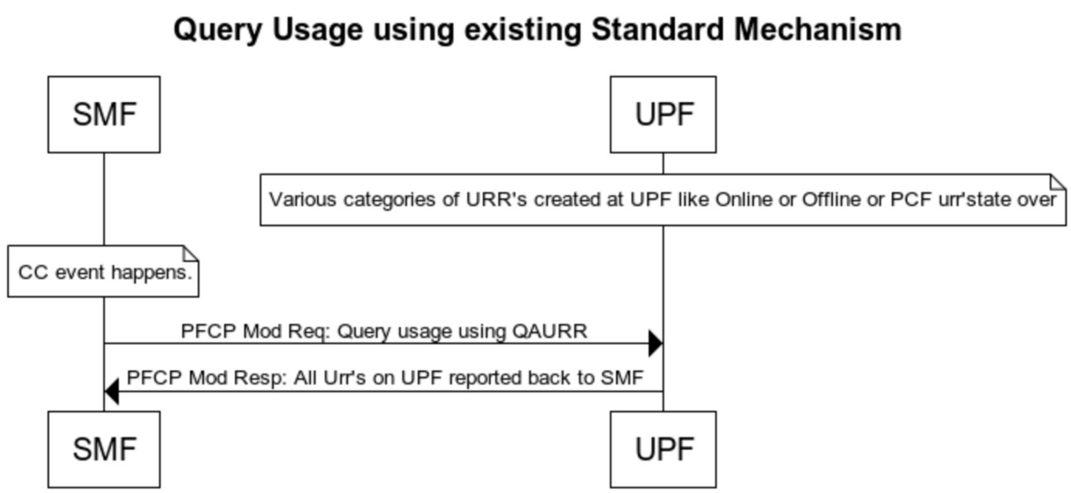

Figure 1: Exemplary Query Usage Employing Standard Mechanisms

The flow that was depicted in Figure 1, above, has a limitation. If a SMF needs to query just a few categories, then a Query All is not possible. If it does not use a Query All, then the SMF needs to query by ID. Such an approach is not optimized and it may result in network congestion. Additionally, the SMF may query only the URR s which are known to the SMF.

A second flow diagram, which is presented in Figure 2, below, depicts elements of a query usage example that employs aspects of the techniques presented herein, specifically a QAURR with a QUERY_INTERFACE IE.

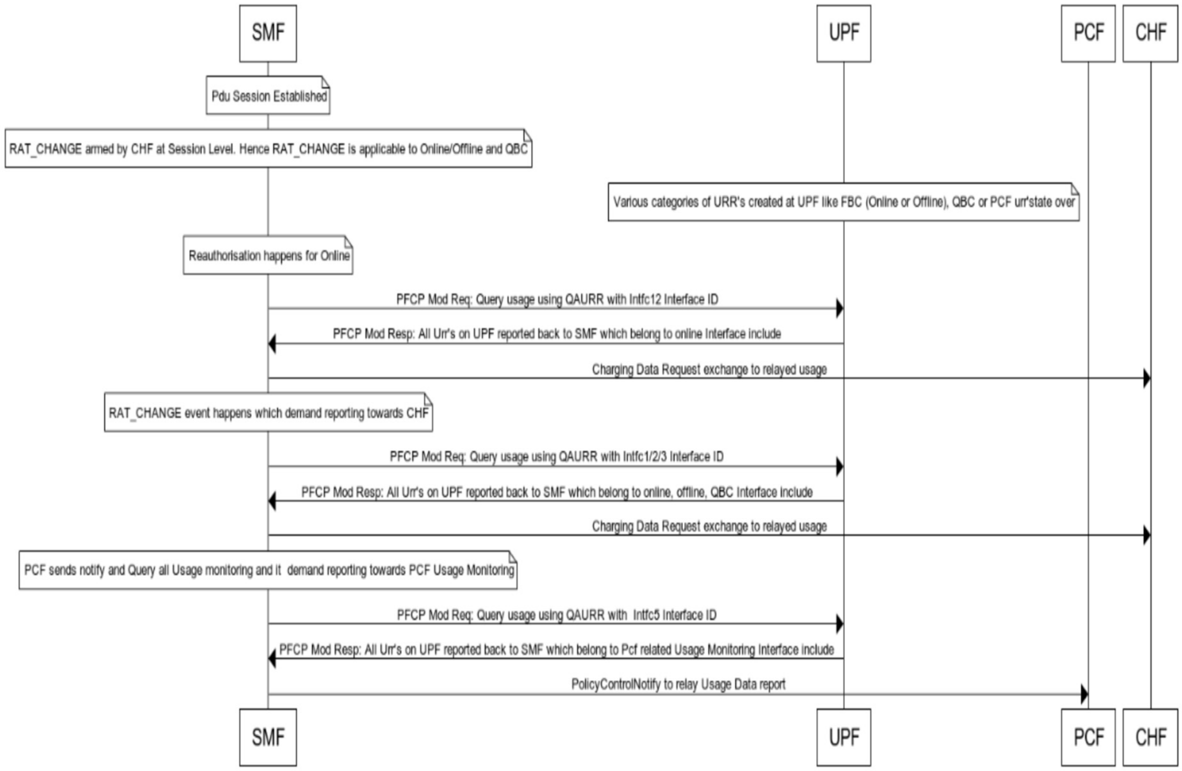

Figure 2: Exemplary Query Usage Employing QAURR

The flow that was depicted in Figure 2, above, has the full flexibility to query usage based on a particular category which needs to be reported. For example, there is no need to query by URR ID if a full category needs to be reported. A SMF can even query what is not known to a CP and which is present on a UP for a particular category.

In summary, techniques have been presented that optimize various of the query usage exchanges that were described in the above narrative. Aspects of the presented techniques support the introduction of a new custom IE, the QUERY_INTERFACE IE on the PFCP interface. Such an IE enables the UP to selectively report multiple URRs which satisfy a category of URR as specified in the IE. Among other things, aspects of the presented techniques provide a mechanism for realizing an optimized PFCP modification request packet by obviating the need to send hundreds of URR identifiers in a request. Additionally, aspects of the presented techniques may also be used within a 3GPP fourth generation (4G) telecommunications environment (e.g., between a control plane Packet Network Data Gateway (PGW-C) and a user plane PGW (PGW-U)).