



US 20150049605A1

(19) **United States**

(12) **Patent Application Publication**
Mann

(10) **Pub. No.: US 2015/0049605 A1**

(43) **Pub. Date: Feb. 19, 2015**

(54) **RULES-BASED OVERLOAD PROTECTION OF A DIAMETER DEVICE**

Publication Classification

(71) Applicant: **ALCATEL-LUCENT CANADA INC.**,
Ottawa (CA)

(51) **Int. Cl.**
H04L 12/801 (2006.01)

(72) Inventor: **Robert A. Mann**, Carp (CA)

(52) **U.S. Cl.**
CPC **H04L 47/12** (2013.01)
USPC **370/230**

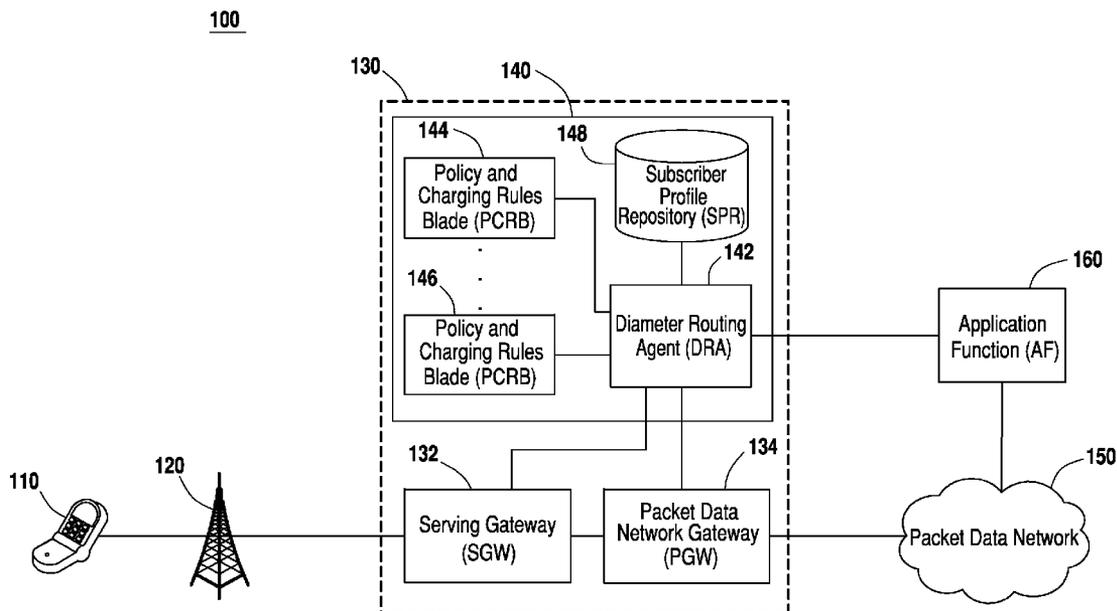
(73) Assignee: **ALCATEL-LUCENT CANADA INC.**,
Ottawa (CA)

(57) **ABSTRACT**

(21) Appl. No.: **13/969,674**

Various exemplary embodiments relate to a method performed by a Diameter Routing Agent (DRA) for processing a Diameter message at the DRA; evaluating an overload state of the DRA; evaluating a shedding rule based upon the overload state of the DRA; and shedding the received Diameter message based upon the evaluation of the shedding rule.

(22) Filed: **Aug. 19, 2013**



100

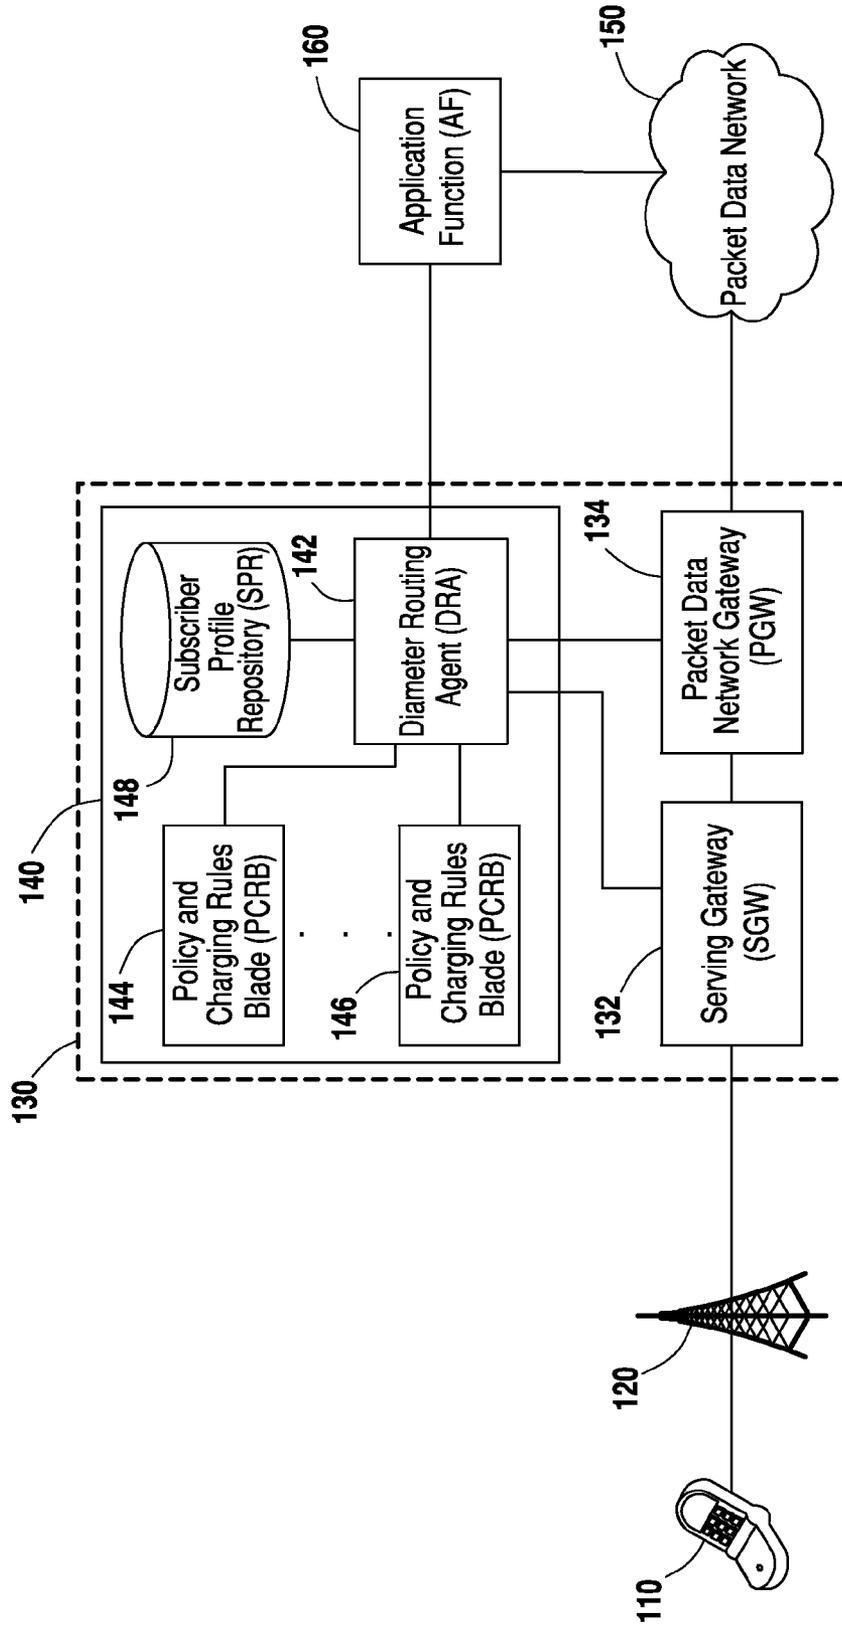


FIG. 1

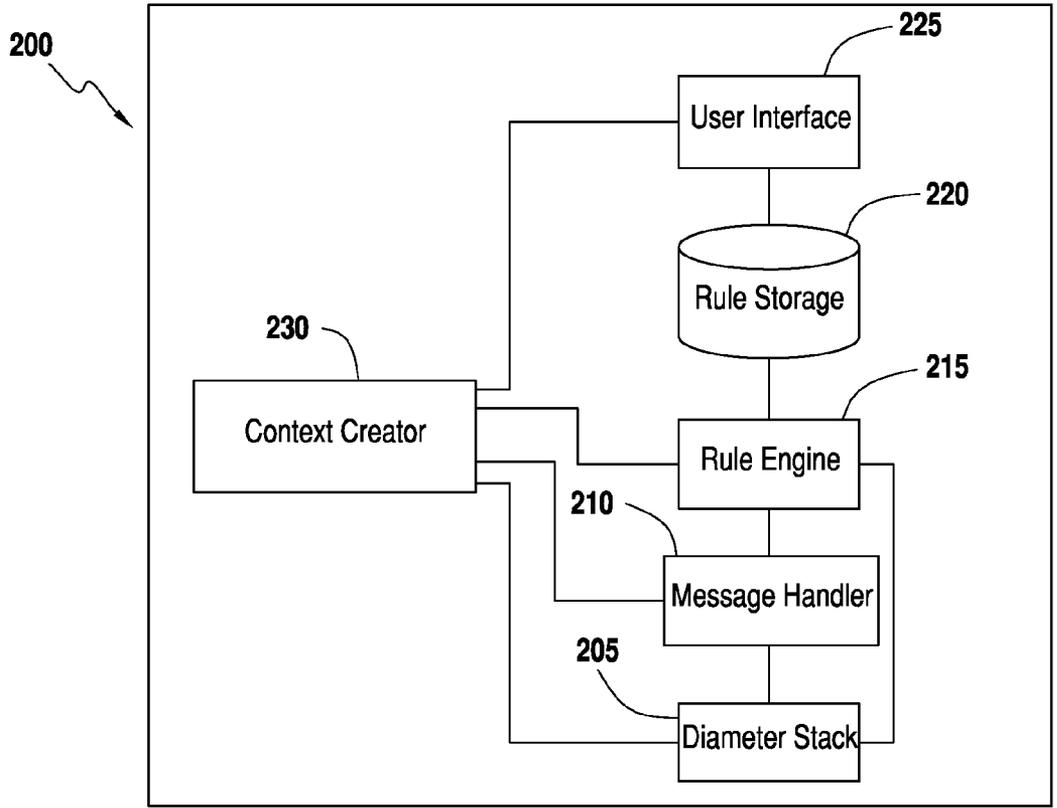


FIG. 2

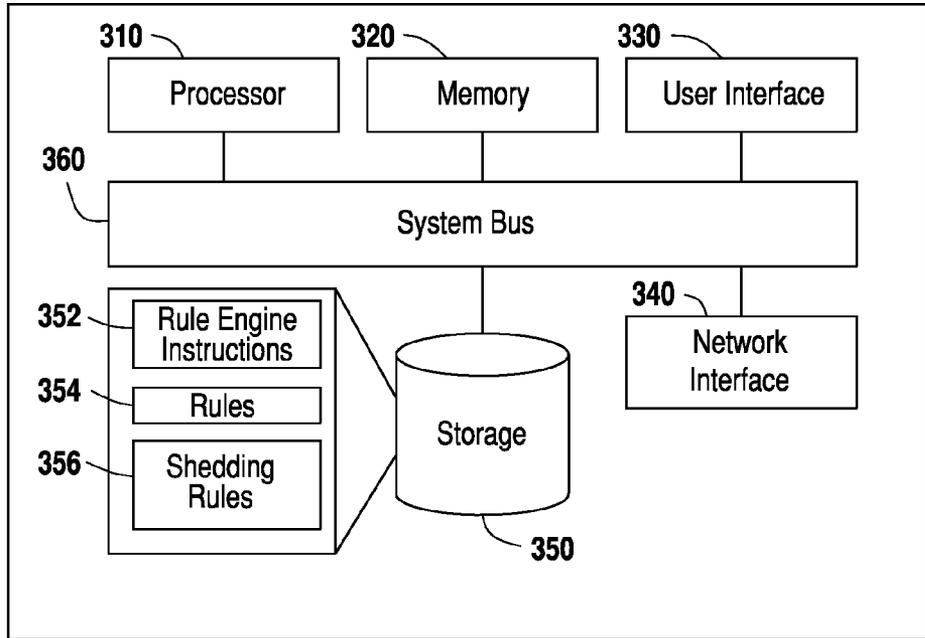


FIG. 3

```
Shed By Level : Rule Table ×
* General Summary
Word Wrap
RULE TABLE: Shed By Level
RULE SYSTEM VERSION: Overload Protection Examples v2 (Draft)
IF [RULE Do nothing if not overloaded]
  (System. Is-Overloaded = 0)
THEN
  (DSC.Return)
ELSE IF [RULE Shed 10 percent at minor]
  (System. Overload-State = 1)
THEN
  (Calculator. Int. Set = 10) AND
  (RULE TABLE: Shed Percentage)
  IF [RULE Shed Percentage]
    // Shed a percentage of messages. It is assumed that
    // Calculator.Int0 contains a value between 1 and 100 - the
    // percentage of requests to shed.
    (Calculator. Random. Next-Int-In-Range(Range = 100) < Calculator. Int0)
  THEN
    (Diameter Answer. Set-Error-Flag = 1) AND
    (Diameter Answer. Result-Code. Set = 3004) AND
    (Calculator. String. Set = DCP in overloaded state. Overloaded components:) AND
    (Calculator. String. Append = System. Overloaded-Components) AND
    (Diameter Answer. Error-Message. Set = Calculator.String0)
ELSE IF [RULE Shed 30 percent at major]
  (System. Overload-State = 2)
THEN
  (Calculator. Int. Set = 30) AND
  (RULE TABLE: Shed Percentage)
  IF [RULE Shed Percentage]
    // Shed a percentage of messages. It is assumed that
    // Calculator.Int0 contains a value between 1 and 100 - the
    // percentage of requests to shed.
    (Calculator. Random. Next-Int-In-Range(Range = 100) < Calculator. Int0)
  THEN
    (Diameter Answer.Set-Error-Flag = 1) AND
    (Diameter Answer. Result-Code. Set = 3004) AND
    (Calculator. String. Set = DCP in overloaded state. Overloaded components:) AND
    (Calculator. String. Append = System. Overloaded-Components) AND
    (Diameter Answer. Error-Message. Set = Calculator. String0)
ELSE IF [RULE Shed 80 percent at critical]
  (System. Overload-State = 3)
THEN
  (Calculator. Int. Set = 80) AND
  (RULE TABLE: Shed Percentage)
  IF [RULE Shed Percentage]
    // Shed a percentage of messages. It is assumed that
    // Calculator.Int0 contains a value between 1 and 100 - the
    // percentage of requests to shed.
    (Calculator. Random. Next-Int-In-Range(Range = 100) < Calculator. Int0)
  THEN
    (Diameter Answer. Set-Error-Flag = 1) AND
    (Diameter Answer. Result-Code. Set = 3004) AND
    (Calculator. String. Set = DCP in overloaded state. Overloaded components:) AND
    (Calculator. String. Append = System.overloaded-Components) AND
    (Diameter Answer. Error-Message. Set = Calculator. String0)
```

FIG. 4

Gx CCR Shedding:Rule T... X

* General
Summary

Word Wrap

RULE TABLE: Gx CCR Shedding

RULE SYSTEM VERSION: Overload Protection Examples v2 (Draft)

RULE SETS:
Gx CCR

IF [RULE Do nothing At Normal]
(System.Is-Overloaded = 0)

THEN
(DSC.Return)

ELSE IF [RULE Shed Gx CCR Establish]
(Gx CCR.CC-Request-Type = 0)

THEN
(RULE TABLE: Shed Gx CCR Establishes By RAT Type
IF [RULE Reject All But EUTRAN Or Emergency]
(Gx CCR.Origin-Realm != emergencyRealm) AND
(Gx CCR.RAT-Type != 1004)
THEN
(Gx CCA.Set-error-Flag = 1) AND
(Gx CCA.Result-Code.Set = 3004))

ELSE IF [RULE Shed Gx CCR Update]
(Gx CCR.CC-Request-Type = 2)

THEN
(RULE TABLE: Shed Gx CCR Updates By RAT Type
IF [RULE Reject All But EUTRAN Updates]
(Gx CCR.RAT-Type = 1004)
// Shed a percentage of messages. It is assumed
// that Calculator.Int0 contains a value between
// 1 and 100 - the percentage of requests to
// shed
(Calculator.Random.Next-Int-In-Range(Range = 100) < Calculator.Int0)
THEN
(Diameter Answer.Set-Error-Flag = 1) AND
(Diameter Answer.Result-Code.Set = 3004) AND
(Calculator.String.Set = DCP in overloaded state. Overloaded components:) AND
(Calculator.String.Append = System.Overloaded-Components) AND
(Diameter Answer.Error-Message.Set = Calculator.String0)))

FIG. 5

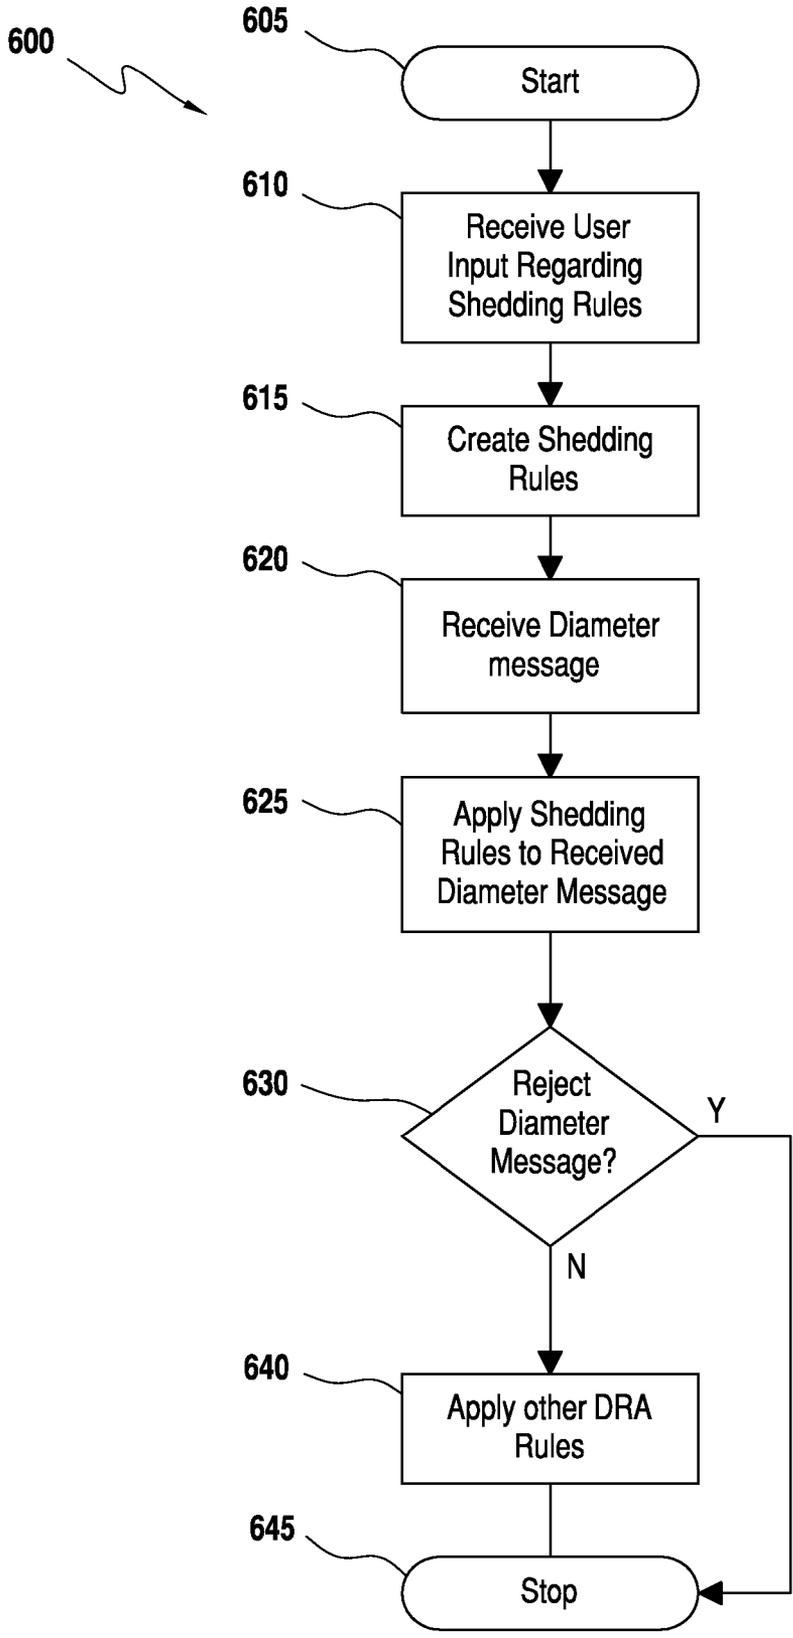


FIG. 6

RULES-BASED OVERLOAD PROTECTION OF A DIAMETER DEVICE

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to the following co-pending applications, which are hereby incorporated by reference for all purposes as if fully set forth herein: application Ser. No. 13/482,690, filed on May 29, 2012, "ORGANIZATION OF DIAMETER ROUTING AGENT RULE SETS;" application Ser. No. 13/482,587, filed on May 29, 2012, "ROUTING DECISION CONTEXT OBJECTS;" application Ser. No. 13/602,579, filed on Sep. 4, 2012, "RULE ENGINE EVALUATION OF CONTEXT OBJECTS;" and application Ser. No. 13/962,071, filed on Aug. 8, 2013, "GENERIC PERSISTENCE IN A DIAMETER ROUTING AGENT."

TECHNICAL FIELD

[0002] Various exemplary embodiments disclosed herein relate generally to communication networks.

BACKGROUND

[0003] Since its proposal in Internet Engineering Task Force (IETF) Request for Comments (RFC) 3588, the Diameter protocol has been increasingly adopted by numerous networked applications. For example, the Third Generation Partnership Project (3GPP) has adopted Diameter for various policy and charging control (PCC), mobility management, and IP multimedia subsystem (IMS) applications. As IP-based networks replace circuit-switched networks, Diameter is even replacing SS7 as the key communications signaling protocol. As networks evolve, Diameter is becoming a widely used protocol among wireless and wireline communications networks.

[0004] One significant aspect of the Diameter protocol is Diameter packet routing. Entities referred to as Diameter routing agents (DRAs) facilitate movement of packets in a network. In various deployments, DRAs may perform elementary functions such as simple routing, proxying, and redirect.

SUMMARY

[0005] A brief summary of various exemplary embodiments is presented below. Some simplifications and omissions may be made in the following summary, which is intended to highlight and introduce some aspects of the various exemplary embodiments, but not to limit the scope of the invention. Detailed descriptions of a preferred exemplary embodiment adequate to allow those of ordinary skill in the art to make and use the inventive concepts will follow in later sections.

[0006] Various exemplary embodiments relate to a method performed by a Diameter Routing Agent (DRA) for processing a Diameter message, the method including: receiving a Diameter message at the DRA; evaluating an overload state of the DRA; evaluating a shedding rule based upon the overload state of the DRA; and shedding the received Diameter message based upon the evaluation of the shedding rule.

[0007] Various exemplary embodiments relate to a Diameter Routing Agent (DRA) for processing a Diameter message, the DRA including: a rule storage configured to store a shedding rule that includes a context object reference; a Diameter stack configured to receive a Diameter message

from an origin device; a rule engine configured to evaluate the shedding rule, wherein the evaluation is based upon an overload state of the DRA; and a message handler configured to shed the received Diameter message based on the evaluation of the rule.

[0008] Various exemplary embodiments relate to a non-transitory machine-readable storage medium encoded with instructions for execution by a Diameter Routing Agent (DRA) for processing a Diameter message, the medium including: instructions for receiving a Diameter message at the DRA; instructions for evaluating an overload state of the DRA; instructions for evaluating a shedding rule based upon the overload state of the DRA; and instructions for shedding the received Diameter message based upon the evaluation of the shedding rule.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] In order to better understand various exemplary embodiments, reference is made to the accompanying drawings, wherein:

[0010] FIG. 1 illustrates an exemplary network environment for a Diameter Routing Agent;

[0011] FIG. 2 illustrates an exemplary Diameter Routing Agent;

[0012] FIG. 3 illustrates an exemplary hardware diagram of a Diameter Routing Agent;

[0013] FIG. 4 illustrates an embodiment of a shedding rule;

[0014] FIG. 5 illustrates another embodiment of a shedding rule; and

[0015] FIG. 6 illustrates an embodiment of a method of applying shedding rules.

[0016] To facilitate understanding, identical reference numerals have been used to designate elements having substantially the same or similar structure or substantially the same or similar function.

DETAILED DESCRIPTION

[0017] The description and drawings merely illustrate the principles of the invention. It will thus be appreciated that those skilled in the art will be able to devise various arrangements that, although not explicitly described or shown herein, embody the principles of the invention and are included within its scope. Furthermore, all examples recited herein are principally intended expressly to be only for pedagogical purposes to aid the reader in understanding the principles of the invention and the concepts contributed by the inventor(s) to furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions. Additionally, the term, "or," as used herein, refers to a non-exclusive or (i.e., and/or), unless otherwise indicated (e.g., "or else" or "or in the alternative"). Also, the various embodiments described herein are not necessarily mutually exclusive, as some embodiments can be combined with one or more other embodiments to form new embodiments. As used herein, the terms "context" and "context object" will be understood to be synonymous, unless otherwise indicated.

[0018] Diameter Routing Agents (DRAs) available today provide only basic functionalities typically defined in hard coding or scripting. As such, users may typically not be empowered to easily and flexibly define more complex behaviors for a DRA. In view of the foregoing, it would be

desirable to provide a method and system that facilitates user definition and extension of DRA message processing behavior.

[0019] FIG. 1 illustrates an exemplary network environment 100 for a Diameter Routing Agent (DRA) 142. Exemplary network environment 100 may be a subscriber network for providing various services. In various embodiments, subscriber network 100 may be a public land mobile network (PLMN). Exemplary subscriber network 100 may be telecommunications network or other network for providing access to various services. Exemplary subscriber network 100 may include user equipment 110, base station 120, evolved packet core (EPC) 130, packet data network 150, and application function (AF) 160.

[0020] User equipment 110 may be a device that communicates with packet data network 150 for providing the end-user with a data service. Such data service may include, for example, voice communication, text messaging, multimedia streaming, and Internet access. More specifically, in various exemplary embodiments, user equipment 110 is a personal or laptop computer, wireless email device, cell phone, tablet, television set-top box, or any other device capable of communicating with other devices via EPC 130.

[0021] Base station 120 may be a device that enables communication between user equipment 110 and EPC 130. For example, base station 120 may be a base transceiver station such as an evolved nodeB (eNodeB) as defined by the relevant 3GPP standards. Thus, base station 120 may be a device that communicates with user equipment 110 via a first medium, such as radio waves, and communicates with EPC 130 via a second medium, such as Ethernet cable. Base station 120 may be in direct communication with EPC 130 or may communicate via a number of intermediate nodes (not shown). In various embodiments, multiple base stations (not shown) may be present to provide mobility to user equipment 110. Note that in various alternative embodiments, user equipment 110 may communicate directly with EPC 130. In such embodiments, base station 120 may not be present.

[0022] Evolved packet core (EPC) 130 may be a device or network of devices that provides user equipment 110 with gateway access to packet data network 140. EPC 130 may further charge a subscriber for use of provided data services and ensure that particular quality of experience (QoE) standards are met. Thus, EPC 130 may be implemented, at least in part, according to the relevant 3GPP standards. EPC 130 may include a serving gateway (SGW) 132, a packet data network gateway (PGW) 134, and a session control device 140.

[0023] Serving gateway (SGW) 132 may be a device that provides gateway access to the EPC 130. SGW 132 may be one of the first devices within the EPC 130 that receives packets sent by user equipment 110. Various embodiments may also include a mobility management entity (MME) (not shown) that receives packets prior to SGW 132. SGW 132 may forward such packets toward PGW 134. SGW 132 may perform a number of functions such as, for example, managing mobility of user equipment 110 between multiple base stations (not shown) and enforcing particular quality of service (QoS) characteristics for each flow being served. In various implementations, such as those implementing the Proxy Mobile IP standard, SGW 132 may include a Bearer Binding and Event Reporting Function (BBERF). In various exemplary embodiments, EPC 130 may include multiple SGWs (not shown) and each SGW may communicate with multiple base stations (not shown).

[0024] Packet data network gateway (PGW) 134 may be a device that provides gateway access to packet data network 140. PGW 134 may be the final device within the EPC 130 that receives packets sent by user equipment 110 toward packet data network 140 via SGW 132. PGW 134 may include a policy and charging enforcement function (PCEF) that enforces policy and charging control (PCC) rules for each service data flow (SDF). Therefore, PGW 134 may be a policy and charging enforcement node (PCEN). PGW 134 may include a number of additional features such as, for example, packet filtering, deep packet inspection, and subscriber charging support. PGW 134 may also be responsible for requesting resource allocation for unknown application services.

[0025] Session control device 140 may be a device that provides various management or other functions within the EPC 130. For example, session control device 140 may provide a Policy and Charging Rules Function (PCRF). In various embodiments, session control device 140 may include an Alcatel Lucent 5780 Dynamic Services Controller (DSC). Session control device 140 may include a DRA 142, a plurality of policy and charging rules blades (PCRBs) 144, 146, and a subscriber profile repository.

[0026] As will be described in greater detail below, DRA 142 may be an intelligent Diameter Routing Agent. As such, DRA 142 may receive, process, and transmit various Diameter messages. DRA 142 may include a number of user-defined rules that govern the behavior of DRA 142 with regard to the various Diameter messages DRA 142 may encounter. Based on such rules, the DRA 142 may operate as a relay agent, proxy agent, or redirect agent. For example, DRA 142 may relay received messages to an appropriate recipient device. Such routing may be performed with respect to incoming and outgoing messages, as well as messages that are internal to the session control device.

[0027] Policy and charging rules blades (PCRB) 144, 146 may each be a device or group of devices that receives requests for application services, generates PCC rules, and provides PCC rules to the PGW 134 or other PCENs (not shown). PCRBs 144, 146 may be in communication with AF 160 via an Rx interface. As described in further detail below with respect to AF 160, PCRB 144, 146 may receive an application request in the form of an Authentication and Authorization Request (AAR) from AF 160. Upon receipt of an AAR, PCRB 144, 146 may generate at least one new PCC rule for fulfilling the application request.

[0028] PCRB 144, 146 may also be in communication with SGW 132 and PGW 134 via a Gxx and a Gx interface, respectively. PCRB 144, 146 may receive an application request in the form of a credit control request (CCR) from SGW 132 or PGW 134. As with an AAR, upon receipt of a CCR, PCRB 144, 146 may generate at least one new PCC rule for fulfilling the application request. In various embodiments, the AAR and the CCR may represent two independent application requests to be processed separately, while in other embodiments, the AAR and the CCR may carry information regarding a single application request and PCRB 144, 146 may create at least one PCC rule based on the combination of the AAR and the CCR. In various embodiments, PCRB 144, 146 may be capable of handling both single-message and paired-message application requests.

[0029] Upon creating a new PCC rule or upon request by the PGW 134, PCRB 144, 146 may provide a PCC rule to PGW 134 via the Gx interface. In various embodiments, such as those implementing the proxy mobile IP (PMIP) standard

for example, PCRB 144, 146 may also generate QoS rules. Upon creating a new QoS rule or upon request by the SGW 132, PCRB 144, 146 may provide a QoS rule to SGW 132 via the Gxx interface.

[0030] Subscriber profile repository (SPR) 148 may be a device that stores information related to subscribers to the subscriber network 100. Thus, SPR 148 may include a machine-readable storage medium such as read-only memory (ROM), random-access memory (RAM), magnetic disk storage media, optical storage media, flash-memory devices, and/or similar storage media. SPR 148 may be a component of one of PCRB 144, 146 or may constitute an independent node within EPC 130 or session control device 140. Data stored by SPR 148 may include subscriber information such as identifiers for each subscriber, bandwidth limits, charging parameters, and subscriber priority.

[0031] Packet data network 150 may be any network for providing data communications between user equipment 110 and other devices connected to packet data network 150, such as AF 160. Packet data network 150 may further provide, for example, phone or Internet service to various user devices in communication with packet data network 150.

[0032] Application function (AF) 160 may be a device that provides a known application service to user equipment 110. Thus, AF 160 may be a server or other device that provides, for example, a video streaming or voice communication service to user equipment 110. AF 160 may further be in communication with the PCRB 144, 146 of the EPC 130 via an Rx interface. When AF 160 is to begin providing known application service to user equipment 110, AF 160 may generate an application request message, such as an authentication and authorization request (AAR) according to the Diameter protocol, to notify the PCRB 144, 146 that resources should be allocated for the application service. This application request message may include information such as an identification of the subscriber using the application service, an IP address of the subscriber, an APN for an associated IP-CAN session, or an identification of the particular service data flows that must be established in order to provide the requested service.

[0033] As will be understood, various Diameter applications may be established within subscriber network 100 and supported by DRA 142. For example, an Rx application may be established between AF 160 and each of PCRBs 144, 146. As another example, an Sp application may be established between SPR 148 and each of PCRBs 144, 146. As yet another example, an S9 application may be established between one or more of PCRBs 144, 146 and a remote device implementing another PCRF (not shown). As will be understood, numerous other Diameter applications may be established within subscriber network 100.

[0034] In supporting the various potential Diameter applications, DRA 142 may receive Diameter messages, process the messages, and perform actions based on the processing. For example, DRA 142 may receive a Gx CCR from PGW 134, identify an appropriate PCRB 144, 146 to process the Gx CCR, and forward the Gx CCR to the identified PCRB 144, 146. DRA 142 may also act as a proxy by modifying the subsequent Gx CCA sent by the PCRB 144, 146 to carry an origin-host identification pointing to the DRA 142 instead of the PCRB 144, 146. Additionally or alternatively, DRA 142 may act as a redirect agent or otherwise respond directly to a request message by forming an appropriate answer message and transmitting the answer message to an appropriate requesting device.

[0035] Every device in a Diameter network is at risk of being overwhelmed by requests sent to it by other devices in the network. To prevent Diameter devices from crashing when overwhelmed, the Diameter device may need the ability to quickly reject/shed the Diameter requests that it is receiving.

[0036] The PCRB 144, 146 may include an overload protection system that monitors the state of various components within the system and detects when they are overloaded. From the usual state of “normal”, the system may escalate on a timed basis through overload states of “minor”, “major” and “critical”. Additionally, a “resource-critical” state is defined which indicates that the system is at risk of crashing. Details of this overload protection system are described in U.S. patent application Ser. No. 13/354,853 entitled “RESOURCE THRESHOLD OVERLOAD PROTECTION,” which is hereby incorporated by reference for all purposes as if fully set forth herein.

[0037] The 3GPP standards are silent on how to proactively shed Diameter requests other than defining a result code of “Too Busy”. In the PCRB 144, 146, the user may be presented with a GUI form that is preconfigured with expected message types organized by Diameter application plus a scenario (typically the command type). The user may specify at what overload levels the messages of that type should be shed at. Because the PCRB is a server, there is a small number of well defined scenarios to be dealt with, so a preconfigured set of scenarios is reasonable and useful.

[0038] For the DRA 142, the situation is more complicated. Currently, the DRA 142 may support up to 17 different Diameter applications, and over time this list may double or triple in size. Because of this large list of applications and how they are used, it would be very difficult to provide a list of pre-defined scenarios that the user can specify a shedding level for.

[0039] The use of the shedding scenarios as predefined in the PCRB in the DRA 142 limit the ability of the user to fine tune shedding. For example, a user may wish to reject only certain requests of the same message type (e.g., Gx CCR) depending on where the request originated from (e.g., a roaming partner) or based on the network type that carries the session (e.g., 3G vs. LTE). This is not possible with preconfigured scenarios, unless the user is allowed to add scenarios and teach the DRA 142 how to detect these scenarios.

[0040] In the DRA 142, instead of providing a preconfigured setup for shedding Diameter requests, a rules engine may be used to fine tune the shedding of requests to whatever specific requirements that particular a user may have. Specifically shedding rules may be used to shed requests in a wide variety of scenarios.

[0041] When the DRA 142 is in an overloaded state, shedding rules may use three different sources of information. The first source of information is the overload state of the system. The following example descriptions of the overload state are based upon those found in U.S. patent application Ser. No. 13/354,853 entitled “RESOURCE THRESHOLD OVERLOAD PROTECTION.” It is noted that other and additional overload states may be defined. If the system is not overloaded, then there is nothing to be done.

[0042] The overload state of the system may be accessible in rules through the following attributes of the ‘System’ context: System.Is-Overloaded—returns 1 if the system is overloaded, 0 if it is not; System.Overload-State—returns a value indicating the overload state of the system, defined as:

0—normal; 1—minor overload; 2—major overload; 3—critical overload; and 4—resource-critical overload.

[0043] Various components of the system may be overloaded. Information regarding overloaded components are as follows: System.Overloaded-Components—a comma-separated list of the names of component(s) that are currently overloaded. For example, the names of system components that may be overloaded are: “Diameter message count”; “CPU Usage”; “Java Heap Memory Usage”; and “Diameter latency”.

[0044] The second input that the shedding rules may use is the specific Diameter message being processed. Any attribute of the Diameter message may be used as criteria when deciding whether to reject the Diameter message. The attributes may be things such as the peer or IP address that the Diameter message was received from, the application and/or command type, or attributes that correspond to AVPs in the Diameter message.

[0045] The third input may be information from a previous Diameter message for the same session that the current Diameter message applies to. By using a generic bindings context, it is possible to remember information about a message as it is routed by the DRA 142. This information can be retrieved later when processing a subsequent Diameter message related to the same session. Such information that may be stored may include any information relating to the session in addition to information related to Diameter messages.

[0046] FIG. 2 illustrates an exemplary Diameter Routing Agent (DRA) 200. DRA 200 may be a standalone device or a component of another system. For example, DRA 200 may correspond to DRA 142 of exemplary environment 100. In such an embodiment, DRA 142 may support various Diameter applications defined by the 3GPP such as, for example, Gx, Gxx, Rx, or Sp. It will be understood that DRA 200 may be deployed in various alternative embodiments wherein additional or alternative applications are supported. As such, it will be apparent that the methods and systems described herein may be generally applicable to supporting any Diameter applications.

[0047] DRA 200 may include a number of components such as Diameter stack 205, message handler 210, rule engine 215, rule storage 220, user interface 225, and context creator 230.

[0048] Diameter stack 205 may include hardware or executable instructions on a machine-readable storage medium configured to exchange messages with other devices according to the Diameter protocol. Diameter stack 205 may include an interface including hardware or executable instructions encoded on a machine-readable storage medium configured to communicate with other devices. For example, Diameter stack 205 may include an Ethernet or TCP/IP interface. In various embodiments, Diameter stack 205 may include multiple physical ports.

[0049] Diameter stack 205 may also be configured to read and construct messages according to the Diameter protocol. For example, Diameter stack may be configured to read and construct CCR, CCA, AAR, AAA, RAR, and RAA messages. Diameter stack 205 may provide an application programmer’s interface (API) such that other components of DRA 200 may invoke functionality of Diameter stack. For example, rule engine 215 may be able to utilize the API to read an attribute-value pair (AVP) from a received CCR or to modify an AVP of a new CCA. Various additional functionalities will be apparent from on the following description.

[0050] Message handler 210 may include hardware or executable instructions on a machine-readable storage medium configured to interpret received messages and invoke rule engine 215 as appropriate. In various embodiments, message handler 210 may extract a message type from a message received by Diameter stack 205 and invoke the rule engine using a rule set that is appropriate for the extracted message type. For example, the message type may be defined by the application and command of the received message. After evaluating one or more rules, rule engine 215 may pass back an action to be taken or a message to be sent. Message handler 210 may then transmit one or more messages via Diameter stack 205, as indicated by the rule engine 215.

[0051] Rule engine 215 may include hardware or executable instructions on a machine-readable storage medium configured to process a received message by evaluating one or more rules stored in rule storage 220. As such, rule engine 215 may be a type of processing engine. Rule engine 215 may retrieve one or more rules, evaluate criteria of the rules to determine whether the rules are applicable, and specify one or more result of any applicable rules. For example, rule engine 215 may determine that a rule is applicable when a received Gx CCR includes a destination-host AVP identifying DRA 200. The rule may specify that the destination-host AVP should be changed to identify a PCRB before the message is forwarded. The rule engine 215 may also receive shedding rules that are evaluated and used to shed received Diameter messages when certain criteria as determined by the shedding rules are met. These shedding rules will be described in more detail below.

[0052] Rule storage 220 may be any machine-readable medium capable of storing one or more rules for evaluation by rule engine 215. Shedding rules may be stored in the rule storage. Accordingly, rule storage 220 may include a machine-readable storage medium such as read-only memory (ROM), random-access memory (RAM), magnetic disk storage media, optical storage media, flash-memory devices, and/or similar storage media. In various embodiments, rule storage 220 may store one or more rule sets as a binary decision tree data structure. Various other data structures for storing a rule set will be apparent.

[0053] It will be understood that, while various components are described as being configured to perform functions such as evaluating rules or accessing context objects based on rules, such configurations may not require any rules to be present in rule storage. For example, rule engine 215 may be configured to evaluate a rule including a context object reference even if no such rule is stored in rule storage 220. Thereafter, if a user adds such a rule to rule storage, rule engine 215 may process the rule as described herein. In other words, as used herein, the phrase “configured to” when used with respect to functionality related to rules will be understood to mean that the component is capable of performing the functionality as appropriate, regardless of whether a rule that requests such functionality is actually present.

[0054] User interface 225 may include hardware or executable instructions on a machine-readable storage medium configured to enable communication with a user. As such, user interface 225 may include a network interface (such as a network interface included in Diameter stack 205), a monitor, a keyboard, a mouse, or a touch-sensitive display. User interface 225 may also provide a graphical user interface (GUI) for facilitating user interaction. User interface 225 may enable a user to customize the behavior of DRA 200. For example,

user interface **225** may enable a user to define shedding rules for storage in rule storage **220** and evaluation by rule engine **215**. Various additional methods for a user to customize the behavior of DRA **200** via user interface **225** will be apparent to those of skill in the art.

[0055] According to various embodiments, rule storage **220** may include rules that reference one or more “contexts” or “context objects.” In such embodiments, context creator **230** may include hardware or executable instructions on a machine-readable storage medium configured to instantiate context objects and provide context object metadata to requesting components. Context objects may be instantiated at run time by context creator **230** and may include attributes or actions useful for supporting the rule engine **215** and enabling the user to define complex rules via user interface **225**. For example, context creator **230** may provide context objects representing various Diameter messages, previous routing decisions, or subscriber profiles.

[0056] Upon DRA **200** receiving a Diameter message to be processed, message handler **210** may send an indication to context creator **230** that the appropriate context objects are to be instantiated. Context creator **230** may then instantiate such context objects. In some embodiments, context creator **230** may instantiate all known context objects or may only instantiate those context objects actually used by the rule set to be applied by rule storage **220**. In other embodiments, context creator **230** may not instantiate a context object until it is actually requested by the rule engine **215**.

[0057] Context creator **230** may additionally facilitate rule creation by providing context metadata to user interface **225**. In various embodiments, context creator **230** may indicate to user interface **225** which context objects may be available for a rule set being modified and what attributes or actions each context object may possess. Using this information, user interface **225** may present a point-and-click interface for creating complex rules. For example, user interface **225** may enable the user to select a desired attribute or action of a context object from a list for inclusion in a rule under construction or modification. Application Ser. No. 13/482,587 and application Ser. No. 13/602,579 provide additional detail regarding the evaluation of context rules and the creation of context objects.

[0058] The contents of Diameter messages may vary depending on the application and command type. For example, an RX RAA message may include different data from a GX CCR message. Such differences may be defined by various standards governing the relevant Diameter applications. Further, some vendors may include proprietary or otherwise non-standard definitions of various messages.

[0059] Upon receiving a Diameter message, message handler **210** passes information regarding the Diameter message to the rule engine **215**. The rule engine **215** may apply various types of rules to the received Diameter message. For example, the rule engine may apply shedding rules, context rules, and routing rules. The rule engine **215** may apply shedding rules first, because if a Diameter message is to be rejected due to the shedding rules and existing DRA **142** loading, such rejection may occur before using resources to apply other rules.

[0060] The rule engine **215** may have a plurality of installed shedding rules to apply to each received Diameter message. In one embodiment, all installed shedding rules may be processed. In another embodiment, shedding rules may be processed until an evaluated shedding rule indicates that the Diameter message is to be rejected. As further evaluation of

the shedding rules is now moot, the rule engine **215** may then begin the evaluation of other types of rules such as context or routing rules.

[0061] The shedding rules may be implemented using a pseudo-code representation of the rule. Such an implementation provides a system user great flexibility in defining rules to meet various requirements and conditions. These shedding rules may use the information described above to evaluate the overload status of the system and to reject Diameter messages accordingly. The shedding rules may be as simple or complex as needed in order to satisfy the specific needs of the user.

[0062] FIG. 3 illustrates an exemplary hardware diagram of a Diameter Routing Agent **300**. The exemplary DRA **300** may correspond to the DRA **200** of FIG. 2 or the DRA **142** of FIG. 1. As shown, the hardware device **300** may include a processor **310**, memory **320**, user interface **330**, network interface **340**, and storage **350** interconnected via one or more system buses **360**. It will be understood that FIG. 3 constitutes, in some respects, an abstraction and that the actual organization of the components of the DRA **300** may be more complex than illustrated.

[0063] The processor **310** may be any hardware device capable of executing instructions stored in memory **320** or storage **350**. As such, the processor may include a microprocessor, field programmable gate array (FPGA), application-specific integrated circuit (ASIC), or other similar devices.

[0064] The memory **320** may include various memories such as, for example L1, L2, or L3 cache or system memory. As such, the memory **320** may include static random access memory (SRAM), dynamic RAM (DRAM), flash memory, read only memory (ROM), or other similar memory devices.

[0065] The user interface **330** may include one or more devices for enabling communication with a user such as an administrator. For example, the user interface **330** may include a display, a mouse, and a keyboard for receiving user commands.

[0066] The network interface **340** may include one or more devices for enabling communication with other hardware devices. For example, the network interface **340** may include a network interface card (NIC) configured to communicate according to the Ethernet protocol. Additionally, the network interface **340** may implement a TCP/IP stack for communication according to the TCP/IP protocols. Various alternative or additional hardware or configurations for the network interface **340** will be apparent.

[0067] The storage **350** may include one or more machine-readable storage media such as read-only memory (ROM), random-access memory (RAM), magnetic disk storage media, optical storage media, flash-memory devices, or similar storage media. In various embodiments, the storage **350** may store instructions for execution by the processor **310** or data upon which the processor **310** may operate. For example, the storage **350** may store rule engine instructions **352** and rules **354** read and acted upon by the rule engine. The storage **350** may also store shedding rules **356** for use in rejecting Diameter messages when the DRA becomes overloaded. It will be apparent that various information described as stored in the storage **350** may be additionally or alternatively stored in the memory **320**. For example, the shedding rules **356** may be additionally, alternatively, or partially stored in memory **320**. Various other arrangements will be apparent.

[0068] FIG. 4 illustrates an embodiment of a shedding rule. The shedding rule in FIG. 4 may shed Diameter messages based upon the overload state. For example, for overload state

1, which indicates a minor overload, the DRA may shed 10% of all Diameter messages received. Further, for overload state 2, which indicates a major overload, the DRA may shed 30% of all Diameter messages received. Further, for overload state 3, which indicates a critical overload, the DRA may shed 80% of all Diameter messages received.

[0069] The rule may first determine if the DRA is overloaded at all. If the DRA is not overloaded, the overload state would be 0, which indicates normal operation of the DRA. Therefore, the shedding rule determines if the overload state is equal to 0, and if it is, then the shedding rule exits as no shedding is needed. If the DRA is overloaded, then a plurality of ELSE IF statements determines whether overload state is 1, 2, or 3. Once the rule identifies the overload state, it may define the percentage of Diameter messages to be shed and generate a random number between 0 and 100. If the random number is less than the shedding percentage associated with the overload state, then the Diameter message is rejected. If the message is rejected, then the Set-Error-Flag may be set to indicate an error, the Result-Code.Set may be set to value that indicates the rejection of the message, and Error-Message.Set may be set to a string indicating that the system is overloaded. In this example, other percentages may be used with each overload state as set by the system user or administrator.

[0070] FIG. 5 illustrates another embodiment of a shedding rule. The shedding rule in FIG. 5 may shed Diameter certain Gx CCR messages. For example, the rule rejects all Gx CCR Establish requests except EUTRAN or Emergency requests. Also, the rule may reject all Gx CCR Update requests except EUTRAN requests.

[0071] In more detail, the rule may first determine if the DRA is overloaded at all. If the DRA is not overloaded, the overload state would be 0, which indicates normal operation of the DRA. Therefore, the shedding rule determines if the overload state is equal to 0, and if it is, then the shedding rule exits as no shedding is needed. If the DRA is overloaded, then the shedding rule next determines if the received Diameter message is a Gx CCR Establish message. If so, the shedding rule then determines if the RAT-Type is either EUTRAN or if the Origin-Realm is an emergency realm. If so, then the received Diameter message is accepted, otherwise the Diameter message is rejected.

[0072] If the Diameter message is not a Gx CCR Establish message, the shedding rule next determines if the received Diameter message is a Gx CCR Update message. If so, the shedding rule then determines if the RAT-Type is EUTRAN. If so, then the received Diameter message is accepted, otherwise the Diameter message is rejected.

[0073] Two example rules are described above using system overload information and information from the received Diameter messages. In addition, as described above, information from a previous Diameter message associated with the same session may also be used in the shedding rules. The shedding rules may be as simple or as complex as the system user desires or requires. Further, a number of shedding rules may be created and applied to each Diameter message received. Also, the shedding rules may be implemented such that one shedding rule can call or invoke another shedding rule. These shedding rules may be applied in a specific desired order as well. Such ordering would allow shedding rules that reject Diameter messages most broadly to be evaluated first in order to conserve resources. The shedding rules may also include the capability to end the further evaluation of other shedding rules in order to conserve resources.

[0074] FIG. 6 illustrates an embodiment of a method of applying shedding rules. The method 600 may be performed by the DRA 142. The method 600 begins 605 and then receives user input regarding shedding rules 610. The user may include the complete definition of shedding rules and the information used in those rules. Also, a set of predefined shedding rule templates may be presented to the user, and the user may select a predefined rule template and provide information to be used to fully define the rule. For example, the shedding rule described above with respect to FIG. 4 has shedding percentages associated with different overload states. A user may select a rule template like the shedding rule of FIG. 4 and then provide the percentage of Diameter messages to reject for each overload state. Many other templates may be provided as well.

[0075] Next, based upon the user input, the method creates the shedding rules 615. The DRA then may receive a Diameter message. Next, the DRA may apply the shedding rules to the received diameter message 625. If the shedding rules reject the Diameter message 630, the method ends 645. If the shedding rules do not reject the Diameter message, the DRA may then apply other DRA rules 640. The other DRA rules may include various routing and context rules used to route Diameter messages to the proper location. The same rule engine may be employed to evaluate the shedding rules as well as the other DRA rules. Such a rule engine provides the benefit of using a single rules engine to implement various functionality in the DRA. Then the method may end at 645.

[0076] The above embodiments have been describe using Diameter messages. The DRA may also route other sorts of messages as well. For example, HTML and SS7 messages. These sorts of messages may create overload situations as well and hence the rules bases shedding may be used to shed these sorts of messages as well.

[0077] According to the foregoing, various embodiments enable robust and dynamic handling of various Diameter messages at a DRA. In particular, by including a rules engine capable of shedding Diameter messages when the DRA is overloaded, the DRA may be protected. Various additional benefits will be apparent from the foregoing disclosure.

[0078] It should be apparent from the foregoing description that various exemplary embodiments of the invention may be implemented in hardware or firmware. Furthermore, various exemplary embodiments may be implemented as instructions stored on a machine-readable storage medium, which may be read and executed by at least one processor to perform the operations described in detail herein. A machine-readable storage medium may include any mechanism for storing information in a form readable by a machine, such as a personal or laptop computer, a server, or other computing device. Thus, a tangible and non-transitory machine-readable storage medium may include read-only memory (ROM), random-access memory (RAM), magnetic disk storage media, optical storage media, flash-memory devices, and similar storage media.

[0079] It should be appreciated by those skilled in the art that any block diagrams herein represent conceptual views of illustrative circuitry embodying the principles of the invention. Similarly, it will be appreciated that any flow charts, flow diagrams, state transition diagrams, pseudo code, and the like represent various processes which may be substantially represented in machine readable media and so executed by a computer or processor, whether or not such computer or processor is explicitly shown.

[0080] Although the various exemplary embodiments have been described in detail with particular reference to certain exemplary aspects thereof, it should be understood that the invention is capable of other embodiments and its details are capable of modifications in various obvious respects. As is readily apparent to those skilled in the art, variations and modifications can be effected while remaining within the spirit and scope of the invention. Accordingly, the foregoing disclosure, description, and figures are for illustrative purposes only and do not in any way limit the invention, which is defined only by the claims.

What is claimed is:

1. A method performed by a Diameter Routing Agent (DRA) for processing a Diameter message, the method comprising:

- receiving a Diameter message at the DRA;
- evaluating an overload state of the DRA;
- evaluating a shedding rule based upon the overload state of the DRA; and
- shedding the received Diameter message based upon the evaluation of the shedding rule.

2. The method of claim 1, wherein:

- receiving a Diameter message includes receiving a plurality of Diameter messages;
- wherein the overload state of the DRA is resource critical; and
- wherein shedding the received Diameter message includes shedding all received Diameter messages.

3. The method of claim 1, further comprising, receiving, via a user interface, a definition of the shedding rule, wherein the definition specifies conditions for shedding Diameter requests.

4. The method of claim 1, wherein evaluating the shedding rule is further based upon attribute value pairs in the received Diameter message.

5. The method of claim 1, wherein evaluating the shedding rule is further based upon information received in a previously received Diameter message.

6. The method of claim 1, wherein the shedding rule defines a percentage of Diameter messages to be shed based upon the overload state of the DRA.

7. The method of claim 1, wherein the shedding rule identifies Diameter messages to shed based upon message type.

8. A Diameter Routing Agent (DRA) for processing a Diameter message, the DRA comprising:

- a rule storage configured to store a shedding rule that includes a context object reference;
- a Diameter stack configured to receive a Diameter message from an origin device;
- a rule engine configured to evaluate the shedding rule, wherein the evaluation is based upon an overload state of the DRA; and
- a message handler configured to shed the received Diameter message based on the evaluation of the rule.

9. The DRA of claim 8, wherein:

- receiving a Diameter message includes receiving a plurality of Diameter messages;
- wherein the overload state of the DRA is resource critical; and
- wherein shedding the received Diameter message includes shedding all received Diameter messages.

10. The DRA of claim 8, further comprising, receiving, via a user interface, a definition of the shedding rule, wherein the definition specifies conditions for shedding Diameter requests.

11. The DRA of claim 8, wherein evaluating the shedding rule is further based upon attribute value pairs in the received Diameter message.

12. The DRA of claim 8, wherein evaluating the shedding rule is further based upon information received in a previously received Diameter message.

13. The DRA of claim 8, wherein the shedding rule defines a percentage of Diameter messages to be shed based upon the overload state of the DRA.

14. The DRA of claim 8, wherein the shedding rule identifies Diameter messages to shed based upon message type.

15. A non-transitory machine-readable storage medium encoded with instructions for execution by a Diameter Routing Agent (DRA) for processing a Diameter message, the medium comprising:

- instructions for receiving a Diameter message at the DRA;
- instructions for evaluating an overload state of the DRA;
- instructions for evaluating a shedding rule based upon the overload state of the DRA; and
- instructions for shedding the received Diameter message based upon the evaluation of the shedding rule.

16. The non-transitory machine-readable storage medium of claim 15, wherein:

- instructions for receiving a Diameter message includes receiving a plurality of Diameter messages;
- wherein the overload state of the DRA is resource critical; and
- wherein shedding the received Diameter message includes shedding all received Diameter messages.

17. The non-transitory machine-readable storage medium of claim 15, wherein evaluating the shedding rule is further based upon attribute value pairs in the received Diameter message.

18. The non-transitory machine-readable storage medium of claim 15, wherein evaluating the shedding rule is further based upon information received in a previously received Diameter message.

19. The non-transitory machine-readable storage medium of claim 15, wherein the shedding rule defines a percentage of Diameter messages to be shed based upon the overload state of the DRA.

20. The non-transitory machine-readable storage medium of claim 15, wherein the shedding rule identifies Diameter messages to shed based upon message type.

* * * * *