Pre-Emptive Static and Dynamic Rate Limiting with CPS vDRA

Contents

Introduction

This document describes rate limit options in DRA, a telecom component that routes Diameter messages and manages network traffic.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Cisco Policy Suite (CPS) Diameter Routing Agent (vDRA)
- Diameter Routing Agent Basics and Specifications

Components Used

The information in this document is based on Cisco Policy Suite DRA.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

DRA is a component in telecommunications networks, particularly within the context of Diameter protocol-based networks. The DRA efficiently routes Diameter messages between different network elements, such as policy servers, charging systems, and other Diameter-enabled devices. Rate limiting is a network traffic management technique used to control the amount of traffic to or from a network element. It helps to ensure that network resources are not exhausted, maintains the quality of service, and prevents misuse or abuse of

the network.

Problem

Each component in the network can handle traffic load based on its rated capacity, but in real time there can be scenarios where the traffic generated is more than what system can handle. Some of them are:

- User behaviour - Activities like streaming events or software updates that generate large amounts of data in a short period. Typically sent from Gateway (Gw) towards DRA.
- Network Congestion - In periods of high network usage, congestion can build up, leading to queued data that is then sent in bursts when capacity becomes available.
- Network resilience mechanisms - Rerouting traffic during outages or maintenance, causing temporary spikes. This can affect the traffic flow in mated sites which does not have any network issue.
- Network Element behaviour – In case of overload and congestion, you can start seeing no responses/timeouts from one or more network elements which can cause reconnection contributing to further overload on the system.
- Gateway flushing - Gateway can flush the existing sessions owing to policy changes, topology change or any maintenance or trouble shooting activity. During these scenarios sessions are cleared and you can receive a burst of Gx Credit Control Request (CCR)-T requests.

Solution

DRA can distribute the load among multiple Diameter servers in order to ensure efficient handling of requests and avoid overloading a single server. In case of server failure, the DRA can redirect messages to alternate servers, ensuring high availability and reliability of the network services.

Rate limiting on the DRA, not just protects the DRA but also other entities by ensuring a controlled flow of messages. Key benefits of rate limiting are:

- Service Continuity - Maintaining continuous service availability by ensuring that critical network components do not get overloaded and prevents outages.
- Scalability - Allowing the network to handle varying loads without degradation in performance.
- Compliance with Service Level Agreements (SLAs) - Ensuring that the network meets its SLAs by maintaining consistent performance and reliability levels.

Static Rate Limit at Load Balancer

This is a straight forward approach, where a fixed threshold is set based on the rated capacity of DRA/Packet Gateway (pGW) and other network elements and does not change based on the network conditions or system resources. By capping the rate of incoming requests you have a predictable outcome on the amount of traffic that DRA processes.

Configurations for static rate limit depends on the use case for which it is applied.

Ingress Rate Limit

Scenario: Bursts from pGW

Thresholds specific to pGWs that are susceptible to these traffic bursts are configured. The value must be arrived based on the regular traffic/peak traffic numbers that can be seen during these bursts.

The threshold numbers can be defined specific to each message type in order to ensure that only the burst traffic is throttled, like, only Gx CCR-I and Gx CCR-T requests from a GW must be throttled but Gx CCR-

U traffic or Gy traffic is forwarded as received.

In this case, you can throttle on the ingress side, that is, DRA throttles the message as soon as it receives it since the purpose here is to reject based on the network element from which it is receiving the request and to avoid processing higher number of requests than the DRA can handle.

The throttle behaviour can be to either reject the message with a particular Error-Code and Error-Message or drop it.

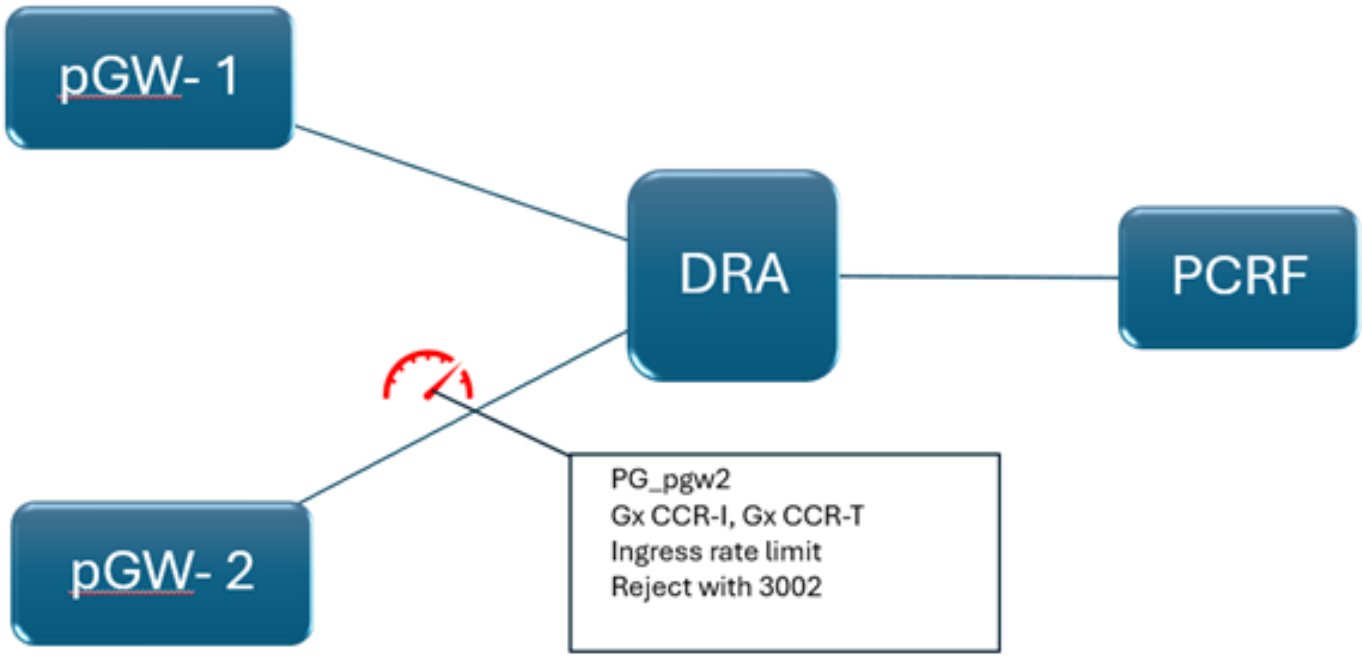

This behaviour can be enabled in CPS vDRA by configuring Custom Reference Data (CRD) Tables 'Peer Rate Limit Profile' and 'Message Rate Limit Profile'. In these CRD tables, you need to configure these values:

Peer Group	A peer group is a logical grouping of Diameter nodes based on their realm and host. You need to configure the peer group that needs to be throttled.
Peer Fully Qualified Domain Name (FQDN)	FQDN (exact or regex match) for the peers in the peer group that you need to rate limit.
Message Direction	Direction of the throttling - Ingress or Egress. In this case - Ingress.
Rate LImit Profile	Message rate limit profile name that used to define the message type that needs to be throttled.
Peer Rate Limit	Rate of requests that are allowed for this Peer Group. This is inclusive of all message types from that Peer Group.

Discard Behaviour	You can choose to drop the request or reject with an Error-Code.
Result Code	Result Code value in case you are rejecting the messages. Not applicable in case the messages are dropped.
Error String	The error string that is used in the response message of the request that are rejected. Not applicable in the case the message is dropped.
Application Identifier	Application ID of the message to be throttled.
Command Code	Command Code of the message to be throttled.
Message/Request Type	Application ID and Request Type of the requests that need to be throttled.
Message Rate Limit	TelePresence Server (TPS) of the request of that message type that is be processed by DRA. Requests beyond this TPS are throttled. This value is per peer in the Peer Group.

Here, is an example, where you are configuring an overall rate limit of 1000 messages for pGW2 and a rate limit of 200 Gx CCR-I and 200 Gx CCR-T. Any request beyond this rate is rejected with a 3002, and an error message indicating that the rate limit is breached.

Message Rate Limit Profile

Filter by: All Visible Columns

CCR_I_T_Limit

Rate Limit Profile Name *	Application Identifier *	Command Code *	Message/Request Type *	Message Rate Limit *	Actions
CCR_I_T_Limit	16777238	272	1	200	
CCR_I_T_Limit	16777238	272	3	200	

Showing 2 out of 2

Show 50 rows 1 out of 1

Egress Rate Limit

Scenario: Protection of Network Element which Receives Request

Consider an example of a Gx transaction where the request is received from pGW and forwarded to the Policy and Charging Rules Function (PCRF). If there are limitations to the amount of data that PCRF can handle, then even though DRA can handle the incoming traffic, you can use DRA to throttle the message at DRA instead of forwarding the request to PCRF and overloading it.

Here you need to throttle on the egress side, that is, DRA throttles the message right before it forwards to the PCRF, based on the PCRF Peer Group that is identified based on DRA routing logic.

The throttle behaviour can be to either reject the message with a particular Error-Code and Error-Message or drop it.

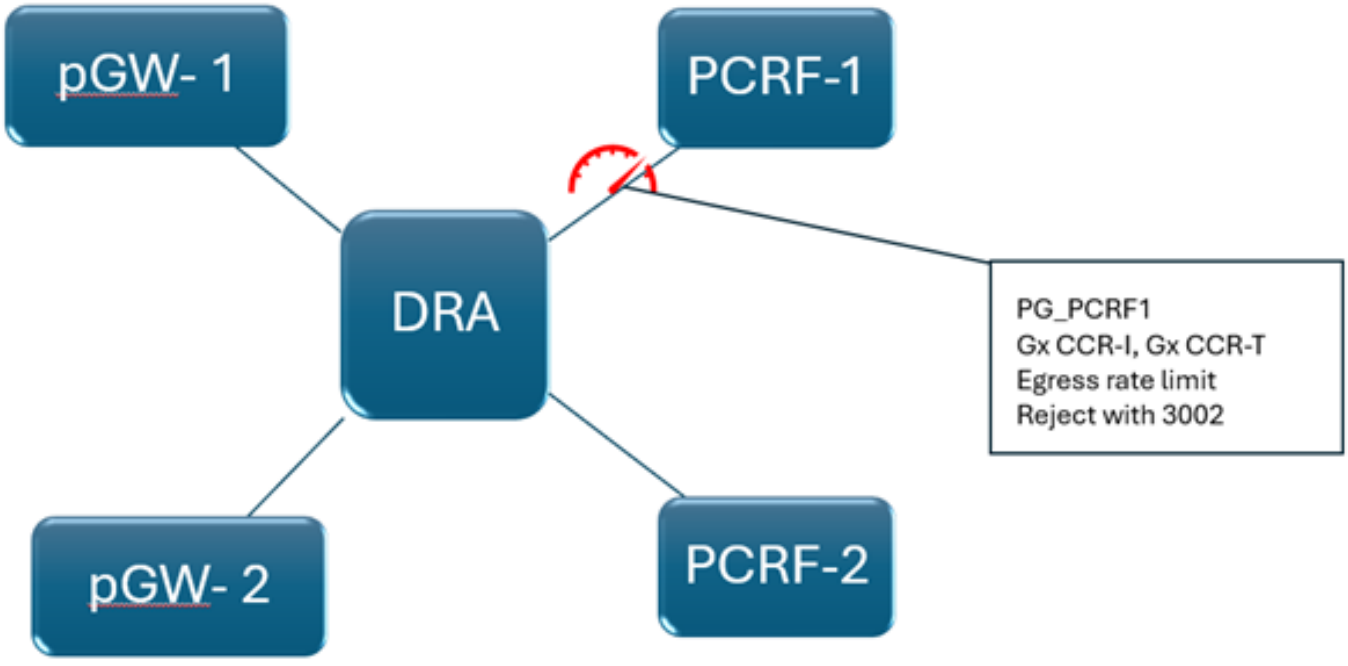

This behaviour can be enabled in CPS vDRA by configuring CRD Tables 'Peer Rate Limit Profile' and 'Message Rate Limit Profile'. In these CRD tables, you need to configure the these values:

Peer Group	A peer group is a logical grouping of Diameter nodes based on their realm and host. You need to configure the peer group that needs to be throttled.
Peer FQDN	FQDN (exact or regex match) for the peers in the peer group that you need to rate limit.
Message Direction	Direction of the throttling - Ingress or Egress. In this case - Egress.
Rate LImit Profile	Message rate limit profile name that used to define the message type that needs to be throttled.
Peer Rate Limit	Rate of requests that is to be allowed for this Peer Group. This is inclusive of all message types from that Peer Group.
Discard Behaviour	You can choose to drop the request or reject with an Error-Code.
Result Code	Result Code value in case you are rejecting the messages. Not applicable in case the messages are dropped.
Error String	The error string that is used in the response message of the request that are rejected. Not applicable in the case the message is dropped.
Application Identifier	Application ID of the message to be throttled.
Command Code	Command Code of the message to be throttled.
Message/Request Type	Application ID and Request Type of the requests that need to be throttled.
Message Rate Limit	TPS of the request of that message type that is processed by DRA. Requests beyond this TPS are throttled. This value is per peer in the Peer Group.

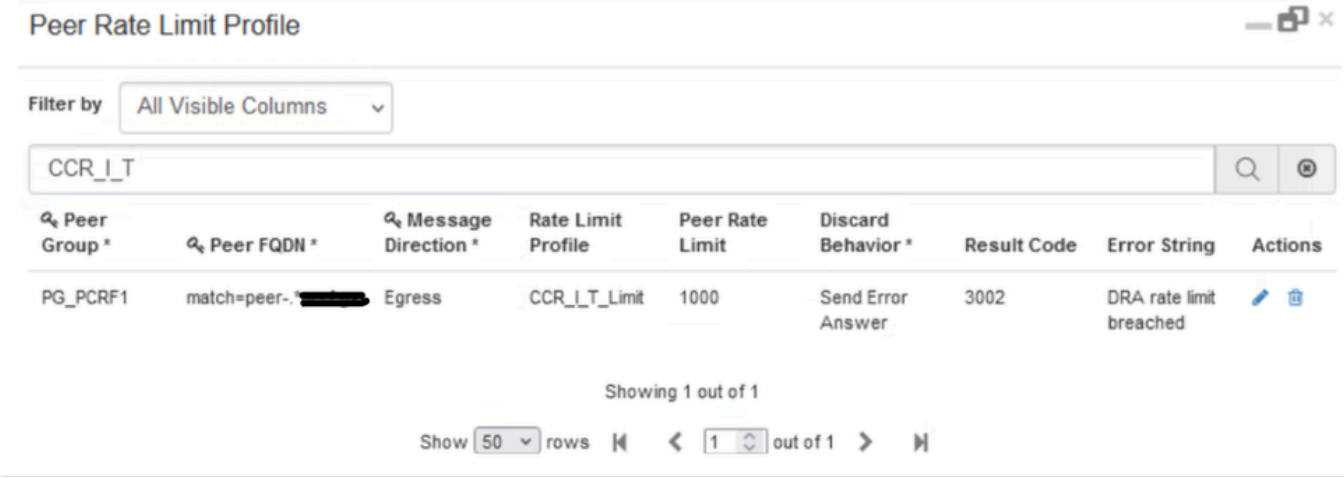

Scenario: Slowness in Network resulting in Traffic Congestion, Causing DRA a Complete/Partial Failure

Consider an example of a Gy transaction that is exchanged between pGW and Online Charging System (OCS). In case of slowness in the network on the DRA-OCS channel (owing to either high traffic from pGW or due to any other network issue), the request gets timed-out due to SLA breach. These timeouts have impact on not just DRA, but the overall network.

DRA resources get held up trying to send the request to OCS over the slow network, resulting in its resources getting exhausted. This results in multiple requests being rejected by DRA, though the rated capacity of DRA is not breached.

This also affects traffic that is not on the DRA-OCS channel. These rejections/time-outs and drops trigger reconnection on multiple network elements.

In this case, you need to throttle on the egress side -DRA throttles the message right before it forwards to the OCS, based on the OCS Peer Group which has capacity limitations or network issues).

The throttle behaviour can be to either reject the message with a particular Error-Code and Error-Message or drop it.

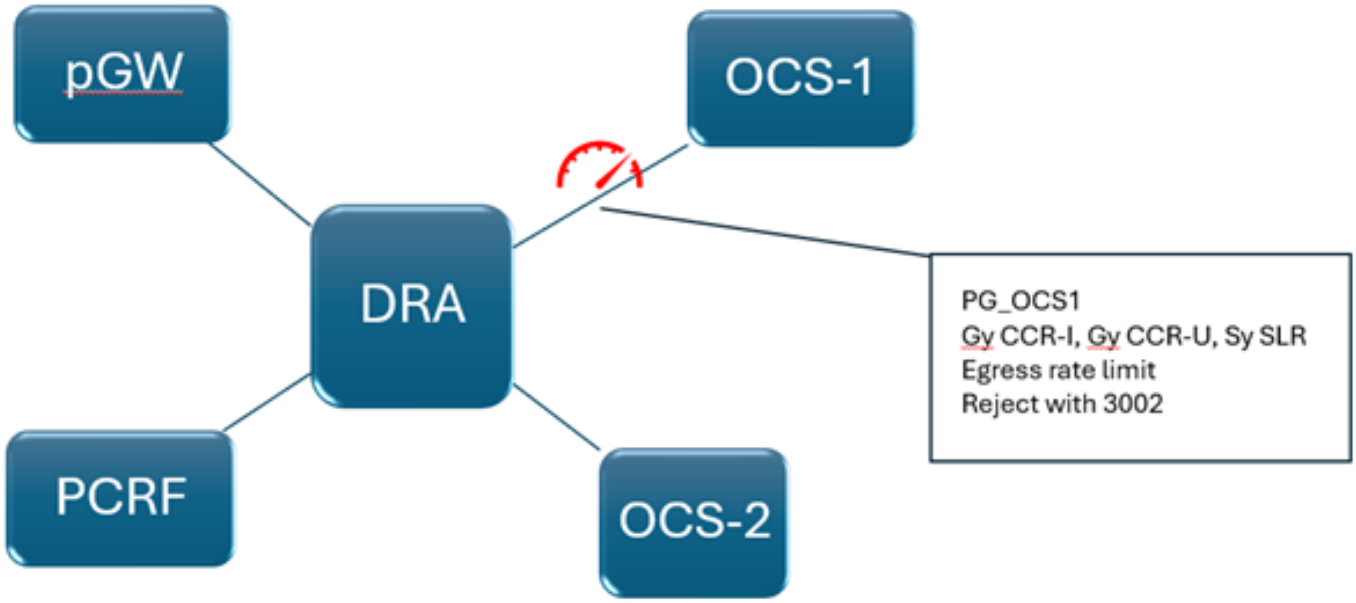

This behaviour can be enabled in CPS vDRA by configuring CRD Tables 'Peer Rate Limit Profile' and 'Message Rate Limit Profile'. In these CRD tables, you need to configure these values:

Peer Group	A peer group is a logical grouping of Diameter nodes based on their realm and host. You need to configure the peer group that needs to be throttled.
Peer FQDN	FQDN (exact or regex match) for the peers in the peer group that you need to rate limit.
Message Direction	Direction of the throttling - Ingress or Egress. In this case - Egress.
Rate Limit Profile	Message rate limit profile name that used to define the message type that needs to be throttled.
Peer Rate Limit	Rate of requests that is allowed for this Peer Group. This is inclusive of all message types from that Peer Group.
Discard Behaviour	You can choose to drop the request or reject with an Error-Code.
Result Code	Result Code value in case you are rejecting the messages. Not applicable in case the messages are dropped.
Error String	The error string that is used in the response message of the request

	that are rejected. Not applicable in the case the message is dropped.
Application Identifier	Application ID of the message to be throttled.
Command Code	Command Code of the message to be throttled.
Message/Request Type	Application ID and Request Type of the requests that need to be throttled.
Message Rate Limit	TPS of the request of that message type that is processed by DRA. Requests beyond this TPS are be throttled. This value is per peer in the Peer Group.

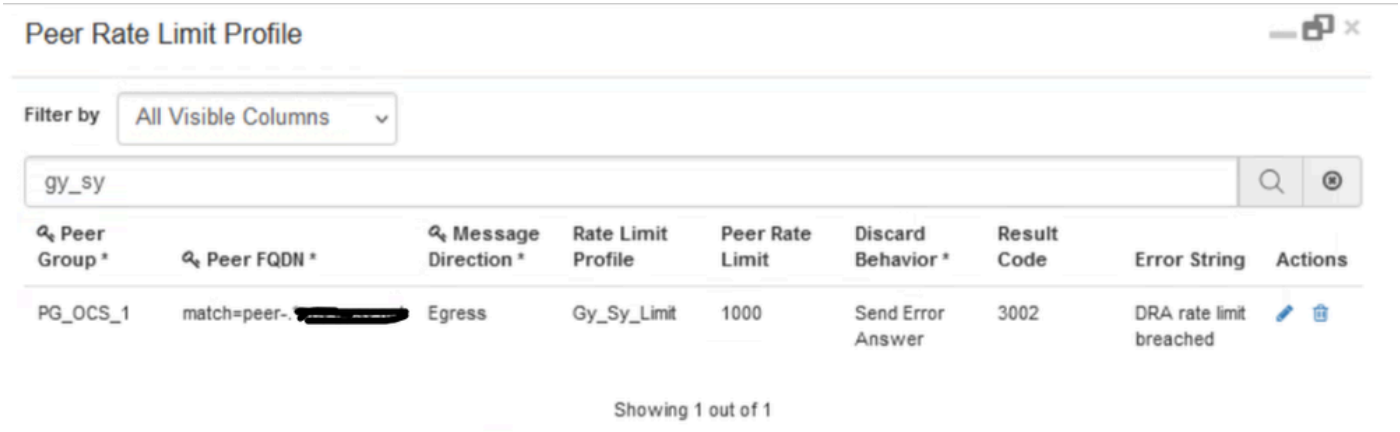

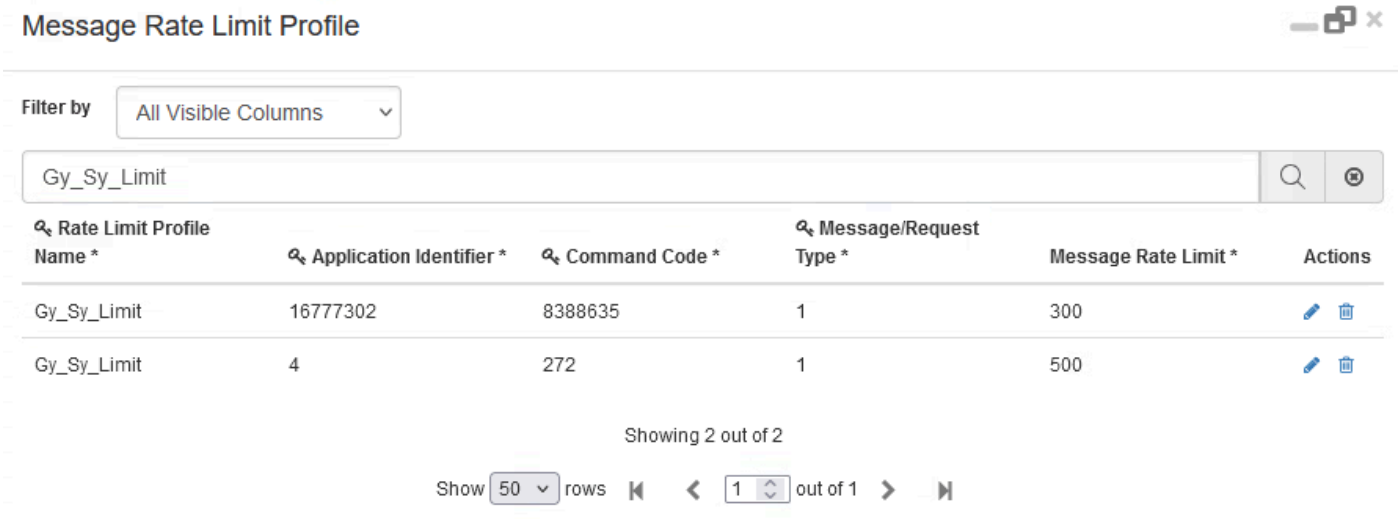

Dynamic Rate Limit

When CCR-I or CCR-T bursts occur, there can be an overload on the database (DB), which can cause the system to destabilize. In order to overcome this, DRA supports dynamic rate limiting (for Gx CCR-I and Gx CCR-T only) based on available DB capacity.

DRA monitors the DB CPU utilization, and whenever the threshold is breached, it throttles the incoming requests. The CPU thresholds for throttling, and the incoming traffic to be throttled are configurable.

Different CPU thresholds with corresponding throttle percentages can be configured. DRA adjusts the throttling level based on the current DB CPU usage. When the CPU usage becomes stable throttling stops gradually.

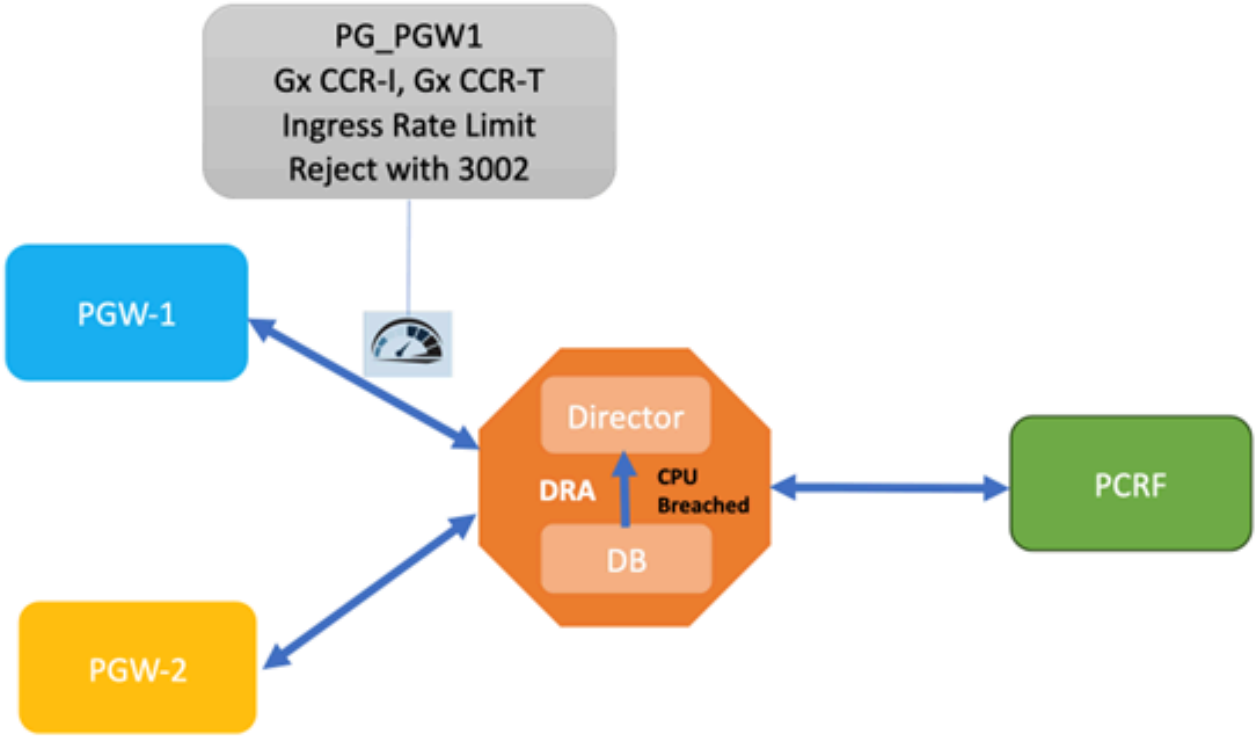

This behaviour can be enabled in CPS vDRA by configuring CRD Tables 'Peer Rate Limit Profile', 'Message Rate Limit Profile', 'Dynamic Peer Rate Limit Profile' and 'Dynamic Throttling DB CPU Profile'. In these CRD tables, you need to configure these values:

Peer Group	A peer group is a logical grouping of Diameter nodes based on their realm and host. In this example you configure the peer group of the pGW.
Peer FQDN	FQDN (exact or regex match) for the peers in the peer group that you need to rate limit.
Message/Request Type	Application ID and Request Type of the requests that need to be throttled. In this example Gx CCR-I, Gx CCR-T.
Message Direction	Direction of the throttling - Ingress or Egress. In this case – Ingress.

Rate Limit Profile	Message rate limit profile name that used to define the message type that needs to be throttled.
Peer Rate Limit	Rate of requests that is allowed for this Peer Group. This is inclusive of all message types from that Peer Group.
Discard Behaviour	You can choose to drop the request or reject with an Error-Code.
Result Code	Result Code value in case you are rejecting the messages. Not applicable in case the messages are dropped.
Error String	The error string that is used in the response message of the request that are rejected. Not applicable in the case the message is dropped.
Application Identifier	Application ID of the message to be throttled.
Command Code	Command Code of the message to be throttled.
Message Rate Limit	TPS of the request of that message type that is processed by DRA. Requests beyond this TPS are throttled. This value is per peer in the Peer Group.
Dynamic Throttling DB CPU Profile	This refers to the CPU profile name, which is used to define the throttle percentage for different CPU range.
DB CPU Utilization Threshold	You can choose the correct value of CPU levels configured as breach limits as per the traffic pattern in your deployment.
Throttle Percentage	The % of ratelimit that is applied when the corresponding CPU level is breached.

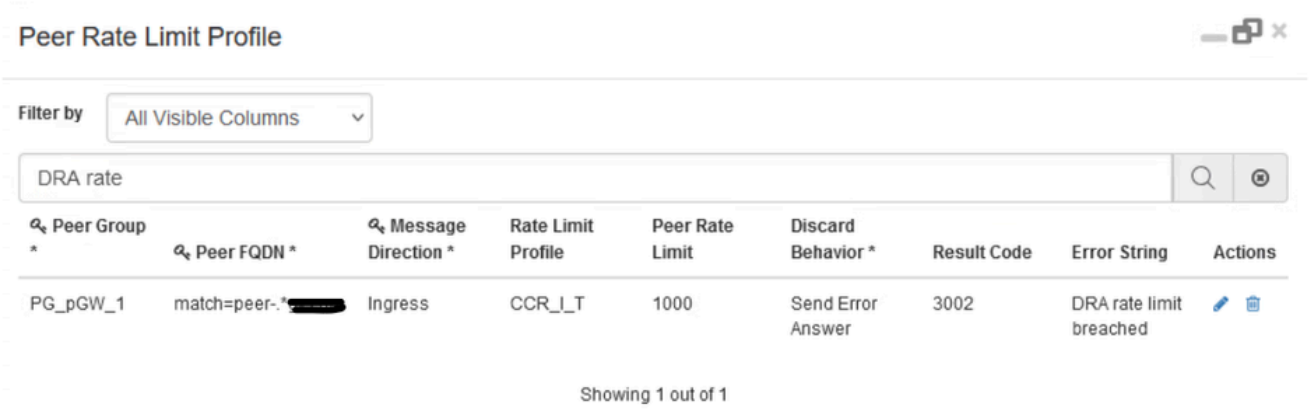

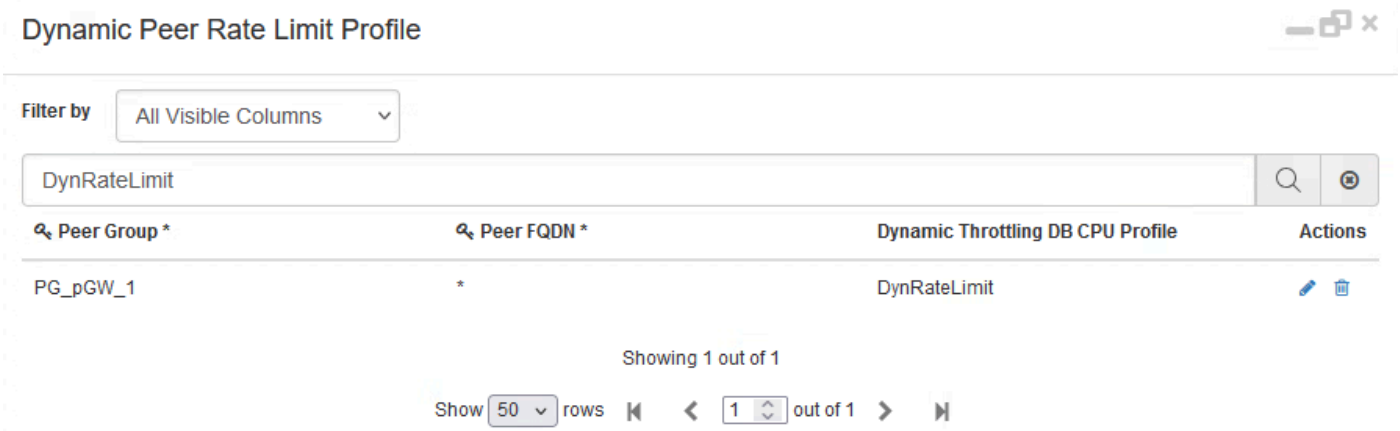

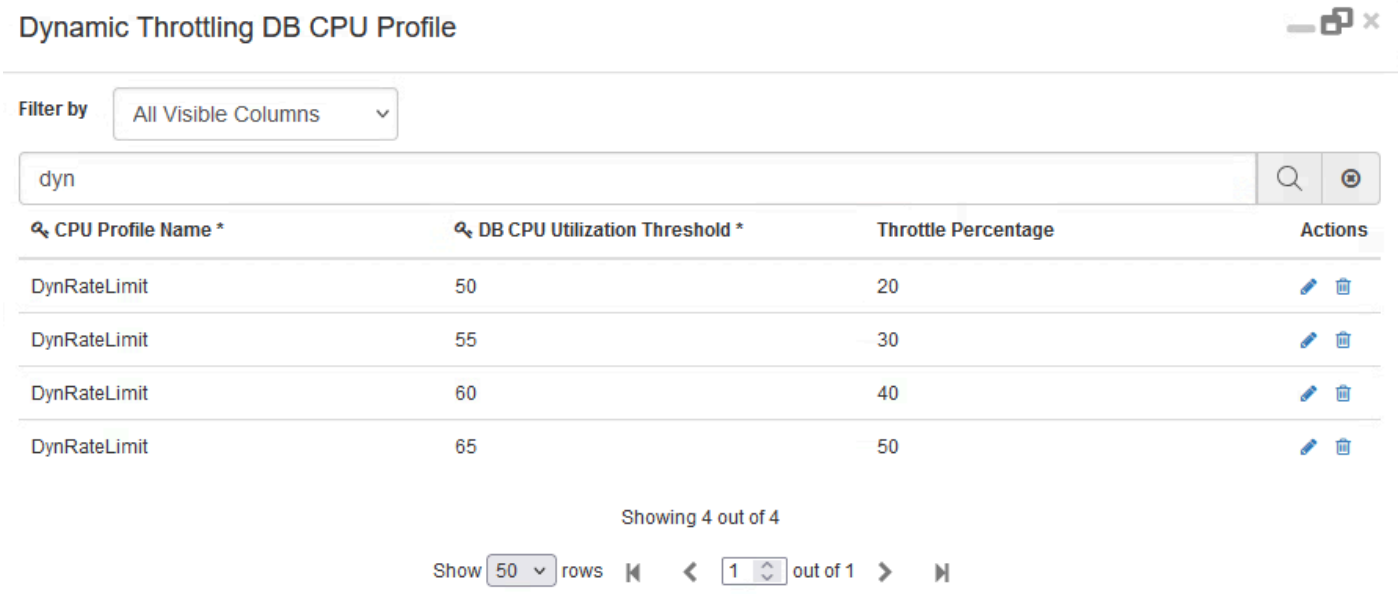

Additionally, this behaviour must be enabled, by choosing the checkbox in Policy Builder under 'DRA Configuration Plugin', under the section 'DRA Dynamic Peer Rate Limiter'.

Reversal Hold Time - Time period for which the CPU utilisation is monitored before reversal is applied.

Reversal Step in % - The percentage of throttling that is reversed.

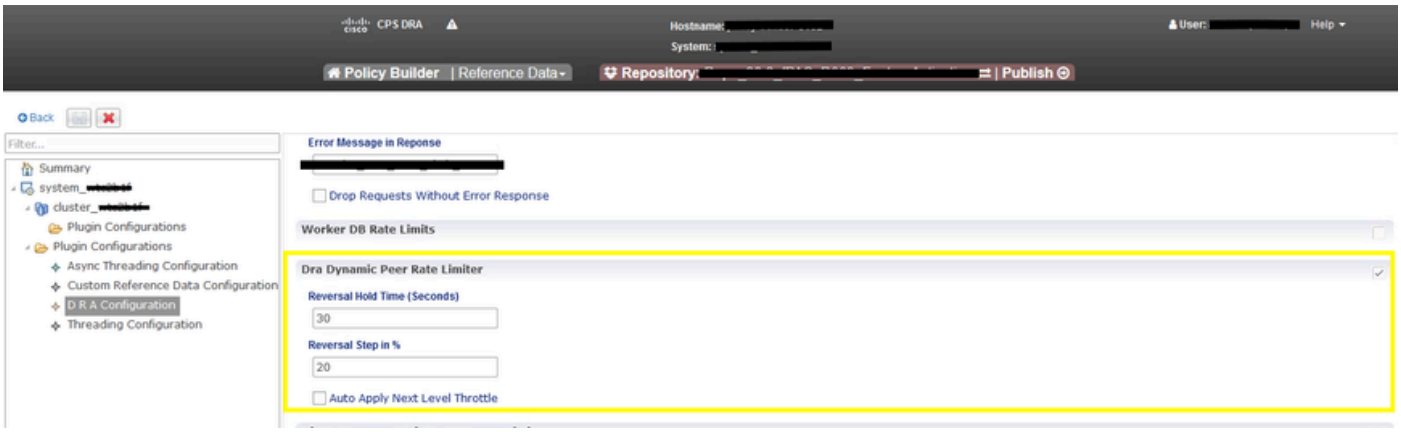

Scenario: Dynamic Rate Limiting based on CPU Utilization

Consider this configuration at DRA:

Static Message Rate Limit: 1000 (it is thus the value of incoming TPS)

Reversal Hold Time: 30s

Reversal step in %: 20%

Whenever DB CPU Utilization crosses the threshold, it refers the 'Dynamic Throttling DB CPU Profile' configuration and throttles the incoming TPS accordingly by notifying the director. Since it is throttling based on ever changing CPU utilization values, you can say its dynamically rate limiting the traffic.

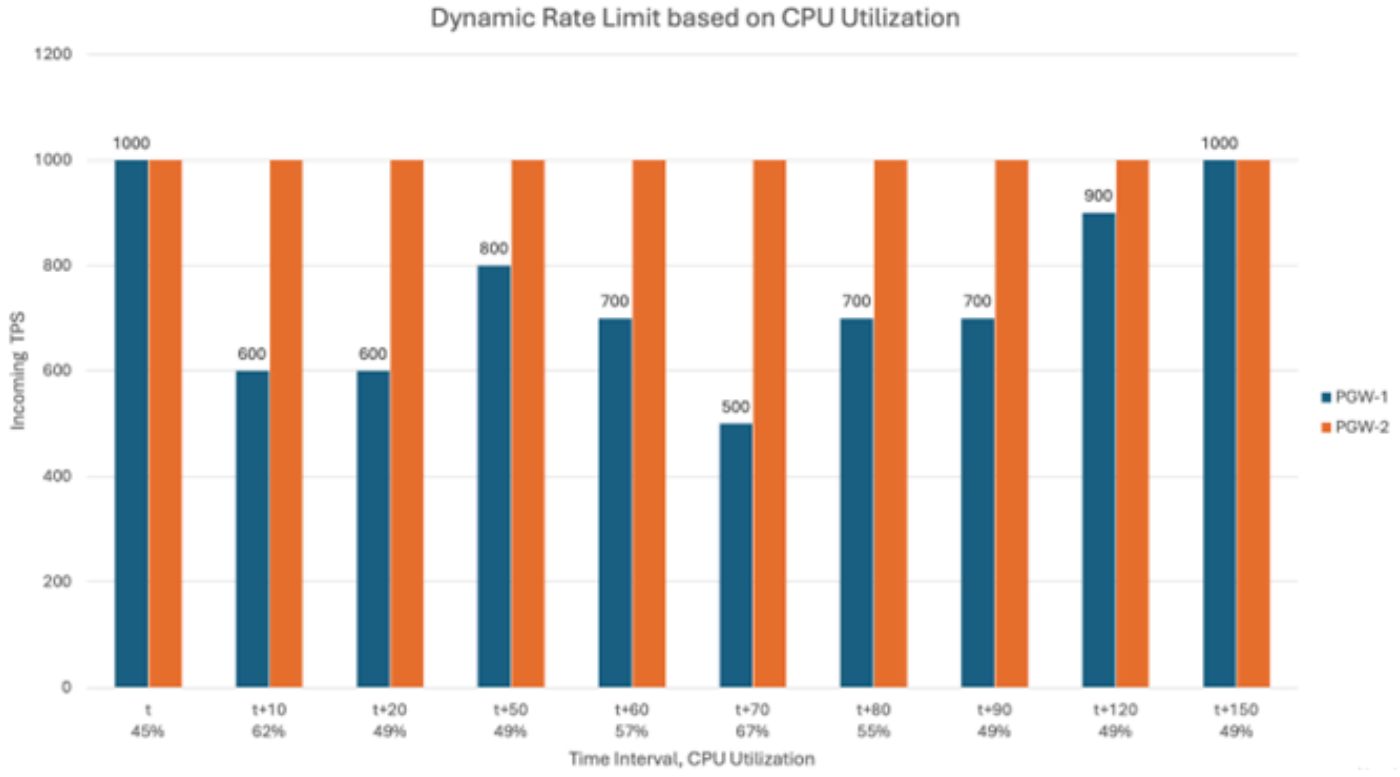

- Initially, DB CPU utilization is under the limit, so no throttling happens. Also, PGW-2 is not having the configuration of Dynamic Rate Limiting and thus no throttling happens there irrespective of CPU

utilization.

- When DB CPU utilization is 62%, traffic is throttled by 40% and effective rate limit is 600 (incoming TPS is 1000, DRA allows only 600).
- If CPU utilization stays between 60-65% then throttling of 40% continues to be applied on configured rate limit of 1000 and effective rate limit is 600 (incoming TPS is 1000, DRA allows only 600.)
- CPU utilization reduces to 49%, reversal of throttling starts at pGW-1.
- If CPU utilization stays at 49% or less for 30 seconds, then throttling is decreased by 20% to 20%. Now effective rate limit is 800 (incoming TPS is 1000, DRA allows only 800.) While reversal, as per configuration, it is done in the steps of 20%.
- When DB CPU utilization increases to 57%, traffic is throttled by 30% and effective rate limit is 700 (incoming TPS is 1000, DRA allows only 700.)
- When DB CPU utilization increases to 67%, traffic is throttled by 50% and effective rate limit is 500 (incoming TPS is 1000, DRA allows only 500.)
- When DB CPU utilization decreases to 55% traffic is throttled by 30% and effective rate limit is 700 (incoming TPS is 1000, DRA allows only 700.)
- If CPU drops to 49% or less for the next 30 seconds interval, then throttling is decreased by 20% to 10% and effective rate limit is 900 (incoming TPS is 1000, DRA allows only 900.)
- If CPU stays furthermore at 49% or less for the next 30 seconds interval, then throttling is decreased by 20% to 0 and there is no rate limit applied as the reversal is complete (incoming TPS is 1000, DRA allows 1000.)