

Why Separate the DRA from PCRF Functionality?



Ian Smith, 2013-27-01

Or Why Not Combine the Functionality if Possible?

To begin, we should ask what are the intended functionalities of a Diameter Routing Agent (DRA) and a Policy Charging & Rules Function (PCRF).

A PCRF is the Policy Decision Point (PDP) virtually cut-and-pasted into IMS and afterwards taken by the 3GPP who awarded it with a new acronym based on the older concept of policy based networking. It is a place where business rules are merged with network actor information (in the case of mobile, subscribers). It acts as a centralized engine for deploying policy to enforcement points in a specific, limited domain (in the case of mobile, usually a MSC, APN, or geographic region). In theory, upstream from the PCRF the business rules themselves are stored in a rules base (the SPR), and the user information is stored in a directory (the HSS or UPSF) and the execution of the policy (rules+user+context) is done downstream by policy enforcement points. However, this theory isn't always strictly translated to practice.

A DRA is a message switching engine that acts on a peer-to-peer network to perform proxy, translation, routing, and relay actions against messages. This is analogous to how the packet switching engine in a network element performs proxy, translation, routing, and relay actions against Ethernet/IP frames (L2-4) but it is performed on the L7 Diameter message in the context of the named interface application.

It is deployed as a strategic point of control to direct Diameter message delivery in the EPC and IMS control plane; facilitates message-proxy, horizontal scaling for control plane elements that cannot scale (load balancing); provides virtual end-points to secure and simplify external communications (virtual servers /DEA); enhances interoperability by performing translation (interop fix-up); bridges between 3GPP and Web 2.0 messaging and architectures; and creates mechanisms for the capture, distribution of, and action on information contained in messages across domains or subsystems in the EPC and IMS (the message switching stratum).

So, given those definitions, should the function of a DRA be separate from the PCRF function?

- Firstly, they don't have anything to do with one another other than they are both Diameter peers. Another way to think of it is to ask why would someone separate the routing and switching function from the SMTP and IMAP function? They both use TCP/IP, so why not put them in the same sheet metal? Deploying a DRA as part of an EPC or IMS functional element is like deploying a L3 switch as the NIC for a server; there isn't anything wrong with doing that, but you'd better have a way to justify the cost and tame the operational complexity of doing that if all you really need is a NIC.
- Secondly, the PCRF is part of the Policy Control and Charging (PCC) subsystem and the Integrated Multimedia Subsystem (IMS). Those are only two functional domains in the EPC and IMS, so by integrating the DRA into the PCRF you've created a hurdle to provisioning that function in the Mobility subsystem, the roaming subsystem/internetworking function, the integration of fixed access, and the visibility and reporting subsystems necessary to actually operate the packet core.
- Thirdly, by choosing an embedded DRA over a message switching stratum for the EPC control plane, the operator is hitching scalability embedded in the DRA function to the expandability of the rest of those functional elements. This is the same conundrum that the dynamic datacenter faces when trying to embed ADC functionality into the elastic compute node; these two functions have different scaling metrics, different utilization curves, and different scaling requirements. Just like OpenFlow and network virtualization doesn't solve that for the cloud, there isn't a viable solution for the elastic packet core.
- Finally, when you don't have a message switching stratum in the EPC control plane (that also extends to the VoLTE functions of the IMS core), you lack a mechanism to apply policy end-to-end on the network, or to deploy offerings that extend policy to roaming customers to enhance value because you don't have a message switching stratum. Instead you have a heterogeneous mix of DRA-like functionalities that may or may not interoperate in a way that can be operationalized, some of which are controlled by the vendor of the function who has a vested interest in not participating in an end-to-end policy architecture.

Put another way, operators who are content to become commoditized dumb pipes shouldn't deploy embedded DRA functionality in their EPC functional elements, because they aren't going to operationalize or monetize their networks with policy based networking driven product offerings. All they really need is a Diameter stack with a static configuration to maximize the fixed configuration economics required for commodity access providers.

Conversely, operators that do intend to pursue value above connectivity should invest in a message switching stratum, because they cannot currently deploy the DevOps apparatus required to tame the operational complexity associated with managing embedded DRA elements in a high capacity dynamic heterogeneous network to deliver policy based offerings that customers will pay extra for (because such an apparatus does not currently exist). Vendors and system integrators will maximize their short-term revenue by selling integrated DRA function at premium rate. And by doing so, they ensure that functional systems are forced to scale in the least cost-effective way as traffic increases. However, they do so with the risk that their solution will be pre-maturely re-evaluated when the need to deploy a message switching stratum arises.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com