An Oracle White Paper
December 2013

The Role of Diameter in All-IP, Service-Oriented Networks

Introduction

Fueled by data-enabled devices (such as smartphones, e-readers, tablets, and netbooks) and compelling applications (such as social networking and mobile video), mobile data traffic is skyrocketing, and there's no slowdown in sight. Indeed, global mobile data traffic has been predicted to increase 26-fold between 2010 and 2015, representing a compound annual growth rate of 92 percent.[1]

Realizing that their 3G networks are not equipped to sustain this high level of traffic growth, today's operators are looking to all-Internet Protocol (IP) networks such as IP Multimedia Subsystem (IMS) and Long Term Evolution (LTE) to provide the bandwidth required to support data-hungry devices and applications and to cost-effectively address the growing gap between traffic and revenue growth.

For years operators have employed Signaling System 7 (SS7) as the international standardized protocol for call setup, mobility management, charging, and intelligent network applications. That's about to change. As operators migrate to all-IP networks, many of the functions performed by SS7 are being replaced by equivalent operations based on the Diameter protocol.

This white paper provides an overview of the Diameter protocol and its many benefits.

[1] "Cisco Visual Networking Index," February 2011.

The Rise of Diameter

Diameter is a message-based signaling protocol designed to provide authentication, authorization, and accounting (AAA) functions in IP-based networks. The name Diameter is a pun on the protocol's predecessor—Remote Authentication Dial-In User Service (RADIUS). Developed in 1998, Diameter was designed to overcome the limitations of RADIUS, and although it isn't directly backward-compatible with that protocol, it does provide an upgrade path.

The Third Generation Partnership Project (3GPP) has adopted Diameter as the primary signaling protocol for AAA, mobility management, and policy and charging control (PCC) in all-IP, service-oriented architectures. The protocol was introduced by the 3GPP in Release 5 for use in IMS networks. Initially, Diameter was deployed in conjunction with other control protocols such as Simple Object Access Protocol (SOAP) and Common Open Policy Service Control (COPS). With Releases 8 and 9, Diameter's use was extended to LTE networks for numerous network functions and interfaces.

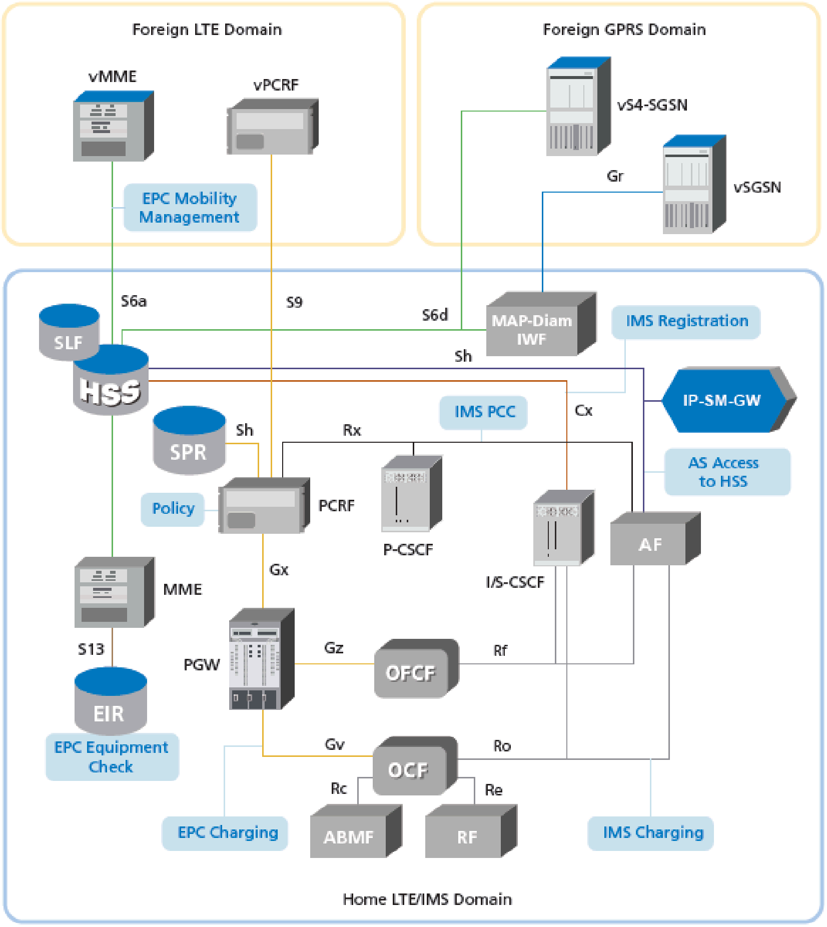

Figure 1. Diameter interfaces in IMS and LTE (Evolved Packet Core [EPC]) networks

Diameter's Structure and Components

The following sections describe the structure and components employed by the message-based signaling protocol to perform its authentication, authorization, and accounting functions.

Messaging

The Diameter message is the basic communication unit used for sending commands and delivering notification between Diameter nodes. Each message must contain a command code that identifies the action to be performed by the message. Because message exchange in Diameter is synchronous, command codes are paired—that is, for every request there is a corresponding answer, which is identified by the same command code. For example, a Diameter node that is requesting authentication information might send an Authentication-Information-Request (AIR). The node that receives the AIR message returns an Authentication-Information-Answer (AIA) message back to the sender.

The Diameter message also contains a set of attribute-value pairs (AVPs), which contain the actual data to be exchanged between Diameter nodes. AVPs include details about AAA as well as routing, security, and capability information. The base Diameter protocol uses some of the AVP values to support required features, such as the following:

- Transport of user authentication information—enabling the Diameter server to authenticate the user
- Transport of service-specific authorization information—to determine whether a user's access request should be granted
- Exchange of resource usage information
- Relaying, proxying, and redirecting of Diameter messages through a server hierarchy

Other AVPs deliver data related to specific applications that employ Diameter. Defined as a base protocol, Diameter provides the minimum requirements needed for AAA, Mobile IPv4, and remote network access applications. It supports basic functions such as reliable transport, message delivery, capabilities negotiation, extensibility, and error handling. The base protocol (defined in RFC 3588) may be used by itself for accounting purposes or in conjunction with other Diameter applications such as those defined by the 3GPP. Each application specifying a service-specific function is created as an extension on top of the Diameter base protocol by addition of new commands or AVPs.

Peer-to-Peer Architecture

Diameter is structured as a peer-to-peer, client/server architecture. A Diameter client typically resides at the edge of the network and is responsible for access control. It receives the request for the user connection and generates messages that request user AAA services. Typically, the Diameter client is a network access server (NAS). The Diameter server performs the actual authentication or authorization of remote users based on profiles. Unlike in the traditional client/server architecture, a Diameter node can act as a server for some requests and as a client in other situations. In addition to clients and servers, the protocol also defines four types of agents: relay, proxy, redirect, and translation.

To route answer messages, every agent must maintain transaction state. When a request is forwarded, its hop-by-hop, locally unique identifier is saved in the message field. When an answer to the request is received, the field reverts to its original value. All nodes in the network are transaction-stateful.

Session-stateful agents keep track of all authorized, active sessions—each of which is bound to a specific service. A session remains active until notification is received that it has ended or until it has expired. Each session has an expiration, which is communicated by Diameter servers with the session timeout AVP. Maintaining session state can be beneficial in several applications, including translating protocols, limiting a user's resources, and per-user or per-transaction auditing. Diameter agents can perform in session-stateful or -stateless mode, depending on the type of request.

Relay Agents

Relay agents forward request messages to other Diameter nodes, based on the information contained in the message—that is, realm (domain), host, and application. Routing decisions are made with the realm routing table, which lists supported realms, known peers, and AVPs. Because relay agents are not involved in policy decisions, they can modify messages only by inserting or removing routing data. They cannot modify any other portion of the message. Relay agents maintain transaction state only.

Proxy Agents

Like relay agents, proxy agents route Diameter messages with the realm routing table. However, unlike relay agents, they can modify the message content and maintain the state of downlink peers such as access devices. These capabilities enable proxies to perform additional functions such as enforcing rules and policies, providing admission control and provisioning, and limiting resource usage. Proxies must maintain transaction state.

Redirect Agents

The redirect agent's sole purpose is to return the routing information needed for Diameter nodes to communicate directly. Upon receipt of a Diameter message, the redirect agent performs a lookup in its routing table. It sends the original sender an answer message that includes the redirect information. After doing so, it drops out of the loop.

The redirect agent acts as a central repository for routing entries, which can be accessed by other Diameter nodes. As a result, the other nodes do not have to maintain the routing data locally, which reduces the overhead associated with updates. Redirect agents do not modify messages and do not maintain session or transaction state.

Translation Agents

Translation agents are proxy agents that provide translation between two protocols such as Diameter and Mobile Application Part (MAP) or Diameter and RADIUS. Because of their function, translation agents must maintain transaction state and be session-stateful.

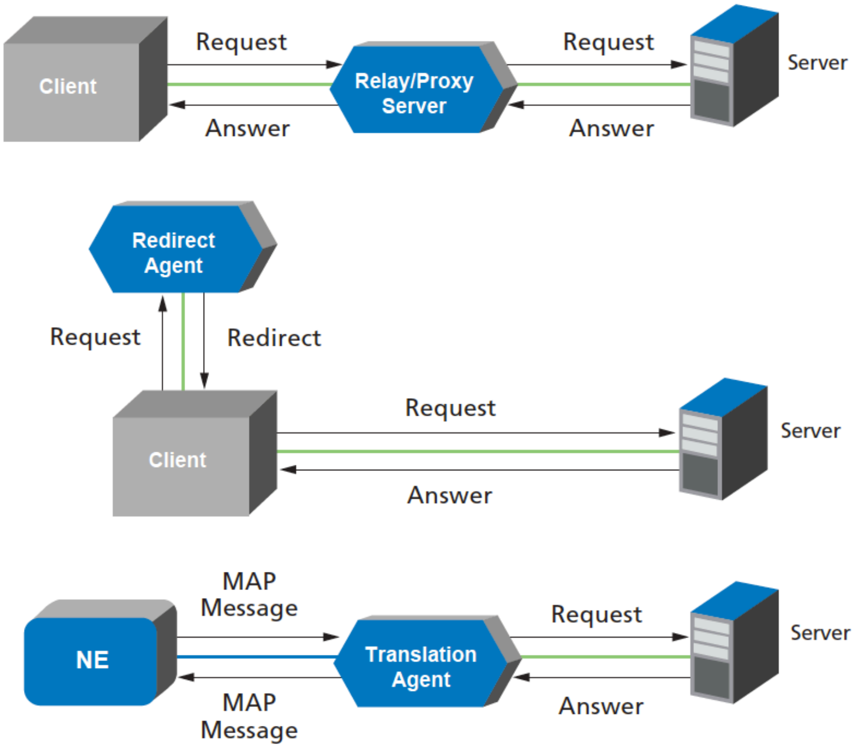

Figure 2. Types of Diameter agents

Peer Discovery

The process by which a node finds another node with which it is going to communicate is called *peer discovery*. Before the advent of the Diameter protocol, the NAS had to be manually configured with the location of its AAA server. In large or complex network deployments, this process was extremely cumbersome. Although all Diameter nodes must support manual discovery, the protocol also provides support for dynamic peer discovery, which reduces the configuration burden on access devices. Once the peer location and routing information are discovered, the Diameter node stores it locally, using two Diameter tables—a peer table and a peer routing table.

Diameter Connections and Sessions

After the peer is discovered, the next stage is to establish a transport connection or a physical link with that peer via the Transmission Control Protocol (TCP) or the Stream Control Transmission Protocol (SCTP). Because Diameter is a peer-to-peer architecture, more than one connection can be established for a node. However, the protocol specifies that a node must establish a connection with at least two peers per realm. The two peers act as the primary and secondary contacts. All messages for a given realm are typically sent to the primary peer. If failover procedures are invoked, all pending requests will be sent to the secondary peer.

When the transport connection is established, the peers must trade capabilities-exchange messages, which detail the peers' identities and capabilities (protocol version number, Diameter applications supported, security mechanisms, and so on). To do so, the initiator sends a Capabilities-Exchange Request (CER) to its peer, which responds with a Capabilities-Exchange Answer (CEA). After the exchange is complete, the peers have the option to negotiate Transport Layer Security (TLS). The connection is then ready to exchange application messages.

Session

A session is a logical connection between two Diameter nodes, which can cross multiple connections. A session refers to the interactions between a client and a server over a given time frame. When a user wants to invoke a service that uses the authentication or authorization portion of an application, the Diameter client sends an authorization request to the local server. The authorization/authentication AVPs are application-specific. As such, they are not defined in the base protocol. The request includes a unique client-generated session ID that identifies that particular session during ongoing communication between the client and the server. During the course of the session, the Diameter server may send a reauthentication or reauthorization request for certain applications such as prepaid service. The server performs this check to avoid further charging once the user has disengaged from the service.

Session termination messages, which can be initiated by clients or servers, are used only for authentication and authorization services and only when the session state is maintained. An accounting stop message is used to terminate an accounting session. When a user session that required authorization by a Diameter server is no longer active, the Diameter client must notify the server by sending it a Session-Termination-Request (STR) message. This notification is critical for tracking purposes as well as notifying stateful agents to release any resources they have provided during the user's session. After the server cleans up the resources associated with the session ID specified in the STR, it returns a Session-Termination-Answer (STA) to the client.

If a Diameter server determines on its end that the session should be terminated, it will send an ABR to the client, which will respond with an ABA.

Benefits of the Diameter Protocol

As networks evolve to all-IP technologies and applications, the need for new AAA mechanisms and requirements has become apparent. The Diameter protocol addresses the shortcomings of its predecessor (RADIUS) and provides the flexibility and extensibility to create the AAA framework for next-generation networks such as IMS and LTE. Its benefits include the following:

- **Reliable transport.** Diameter—which uses either TCP or SCTP—runs over a reliable transport layer with well-defined transport behavior. In contrast, RADIUS employs User Datagram Protocol (UDP), which does not define retransmission behavior and can result in reliability issues across implementations. In addition, UDP doesn't have any means for the receiving node to control the flow of data from the sending node. TCP and SCTP—which are both connection-oriented transport protocols—include mechanisms for controlling flow and adapting to network congestion.

- **Better proxying.** With Diameter, each node in the message path can detect a failure in its next-hop peer. When a proxy detects a failure, it performs the failover locally and automatically retransmits the pending request. With RADIUS, the NAS is responsible for all retransmissions—proxies cannot retransmit requests. Because the NAS does not know whether the failure is remote or local, it may retransmit to an inappropriate next-hop peer.
- **Improved security.** RADIUS has only hop-by-hop security, and the use of the IP Security (IPSec) protocol is optional. AVPs cannot be secured as they pass between the NAS and the home server. Proxy servers can tap into confidential information such as accounting data or modify messages without the endpoint's knowledge. The Diameter protocol improves security by providing both hop-by-hop and end-to-end security. Diameter ensures the integrity and confidentiality of AVPs with digital signatures and encryption.
- **Capability negotiation.** RADIUS clients and servers have no knowledge of each other's capabilities. As a result, they may not be able to successfully negotiate a service; in some cases, the nodes may not even know which service is implemented. In contrast, Diameter clients and servers must exchange their capabilities *before* they can pass messages. This means that each node has full knowledge of the protocol version number, Diameter applications, and security mechanisms supported by its peer before any data is exchanged.
- **Dynamic peer discovery.** With RADIUS, the names and addresses of clients and peers, along with corresponding shared secrets, must be manually configured. In large or complex networks, this approach creates an enormous administrative burden. Diameter enables dynamic peer discovery via Domain Name System (DNS), and transmission-level security supports the dynamic derivation of session keys.
- **Roaming support.** Diameter's explicit support for proxy chaining, auditability, and transmission-level security enables secure, scalable roaming. Because RADIUS lacks these features, it is susceptible to attacks from external parties as well as to roaming partner fraud.

Conclusion

Diameter is emerging as the AAA protocol of choice for all-IP IMS and LTE networks. Virtually all of the industry's standards organizations support the use of the protocol—and with good reason: Diameter provides the scalability, flexibility, reliability, and security to deal with the complex signaling required to make good on the promise of advanced mobile data services and applications. The protocol's extensibility ensures its ability to handle not just existing access technologies and use cases but also those that will emerge in the future.

Oracle is committed to developing practices and products that help protect the environment

The Role of Diameter in All-IP,
Service-Oriented Networks
December 2013

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Hardware and Software, Engineered to Work Together