

New VoLTE Policy Management Architecture: Improve Subscriber Experience and Lower Cost of Growth

What You Will Learn

This white paper explains the unique challenges of policy management for Voice over LTE (VoLTE) and how the Cisco Policy solution solves them. This solution:

- Scales cost-effectively and easily as your service grows
- Can be deployed across geographic areas to provide high availability and redundancy
- Reduces latency to provide an excellent subscriber experience
- Does not require a Diameter Routing Agent for session correlation, which lowers capital and operational costs

VoLTE: Big Opportunity, New Challenges

Offering VoLTE affects profitability in two ways. One is by creating a great voice experience that helps to attract and retain subscribers. VoLTE provides guaranteed quality of service and fast call-setup time. It also supports new services such as call forwarding, call holding, conferencing, and combined voice, video, and chat. These capabilities help to ward off the threat from over-the-top (OTT) providers.

At the same time, VoLTE can lower network costs. For example, it increases spectrum efficiency. It typically doubles capacity for voice calls, or increases data capacity by 50 to 100 percent. And managing both voice and data on the evolved packet core network, instead of on separate networks, can reduce operations and maintenance costs by 50 percent.

VoLTE requires a policy management solution. Even if you already have a policy management solution for your 4G LTE data services, you will need a new one for VoLTE, to:

- Establish and release the bearer for voice traffic on behalf of the IP multimedia system (IMS) domain
- Forward the bearer allocation status and subscriber location from the packet core to the IMS domain
- Forward charging information from the IMS domain to the evolved packet core (optional)
- Handle supplementary services, such as call forwarding and call holding, that you deliver from the IMS domain

VoLTE Policy Management: Four Requirements

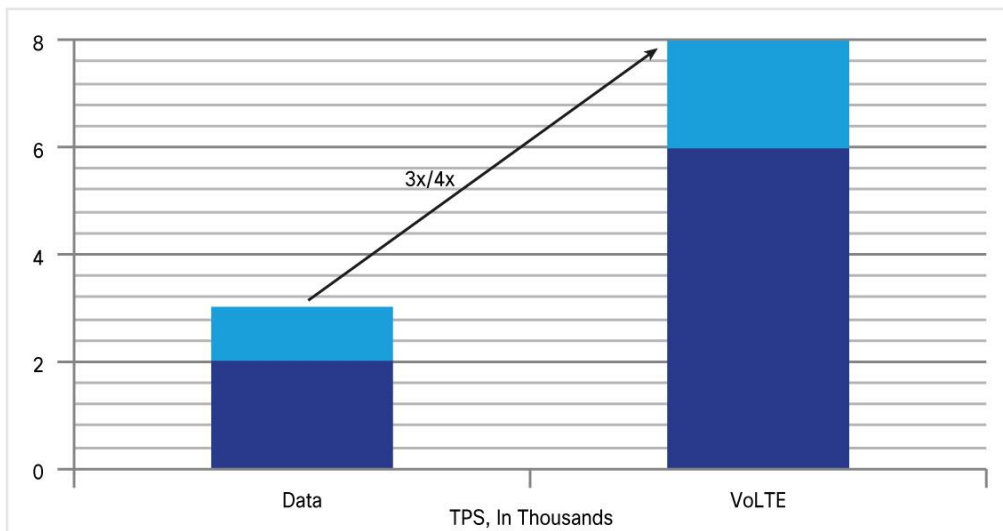
VoLTE changes the game for policy management because it imposes much heavier demands than 4G data services. The policy management solution needs to meet the following four requirements.

Scale to Hundreds of Thousands of Transactions per Second

Adding VoLTE typically triples or quadruples the number of transactions per second (Figure 1). The policy management solution needs to support hundreds of thousands of transactions per second. To keep costs down as volume grows, you need a solution that can scale linearly. This requires:

- Virtualized applications
- Congestion control capabilities to help manage unanticipated spikes
- An **n+1** redundancy architecture, which keeps costs down compared to **1+1** redundancy architectures

Figure 1. Adding VoLTE Can Triple or Quadruple TPS



Latency

For data services, the subscriber experience remains good even with latency of around 100 milliseconds. But for a good experience with VoLTE, latency cannot exceed 15 milliseconds, even during peak times.

Resiliency

Policy management solutions for data servers typically use a standard high-availability architecture. If one server fails, any other server in the data center can take over. If the entire site fails, however, session information is lost. That's a problem for VoLTE. Subscribers who tolerate brief interruptions to data services will not tolerate dropped calls.

Therefore, the policy management architecture for VoLTE should be geographically redundant. Session information is replicated both within the data center and between data centers. An active-active architecture is preferable to an active-standby architecture, because it supports more stringent service-level agreements (SLAs).

Session Correlation

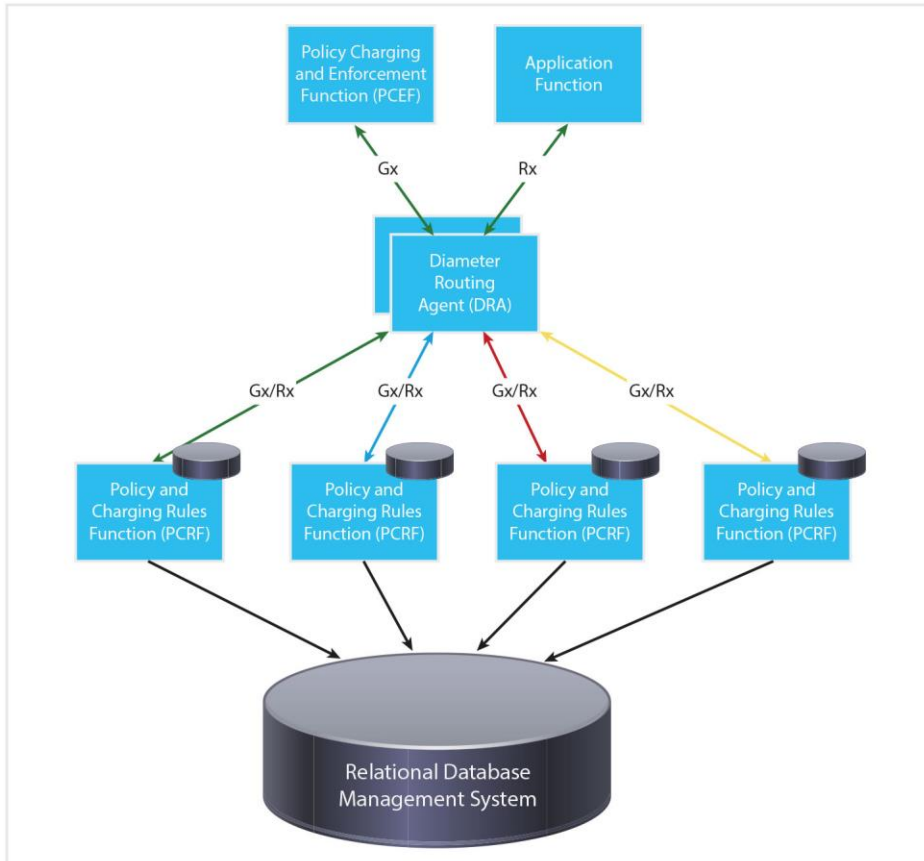
Policy management solutions for data services only need to handle one session, from the policy charging and enforcement function (PCEF). For VoLTE, the policy management solution needs to correlate two sessions: one from the PCEF and another one from the application functions. Traditional policy management solutions provide a very limited session correlation capability, challenged by scalability needs.

Why Traditional Policy Management Architectures Fall Short for VoLTE

Most policy management solutions do not meet all four requirements for VoLTE. Problems include:

- **Poor scalability:** The main reason for poor scalability is that application logic and session persistence are combined in the same node. Another reason is the use of relational database management systems, which do not scale easily and are difficult and expensive to maintain.
- **Latency, which degrades the voice experience:** The DRA that balances requests across the nodes increases latency (Figure 2). It also adds to the number of network elements, increasing capital and operational expense.
- **Lack of support for geographic redundancy**
- **High costs:** In addition to being difficult to scale, relational databases are costly. Costs include servers, software licensing, and the operational expense of constant tuning.

Figure 2. Traditional Policy Management Architectures Require a Diameter Routing Agent, Increasing Latency

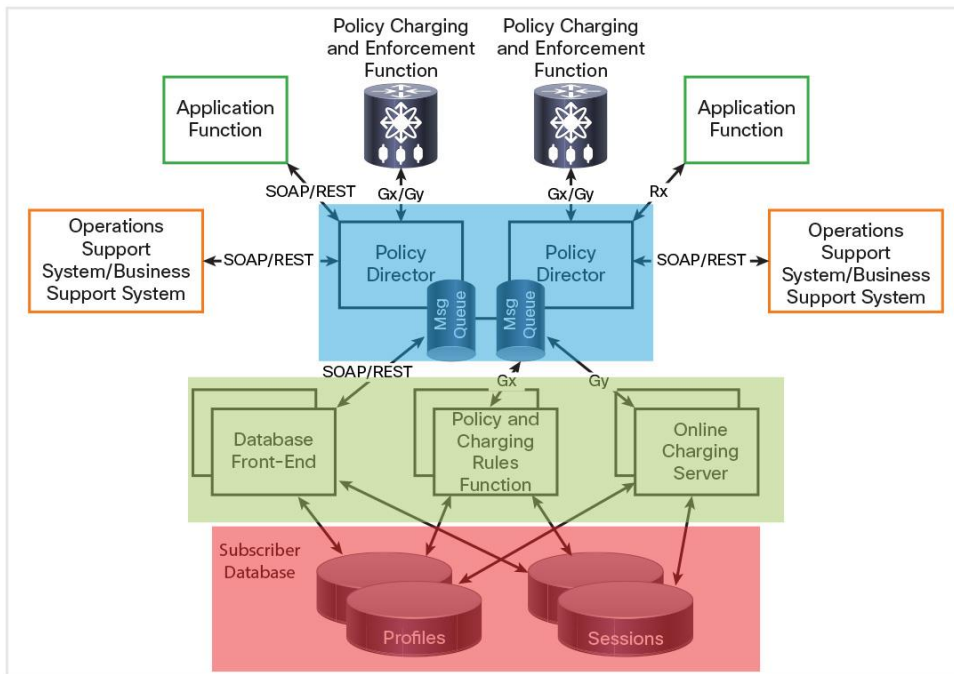


Cisco Policy Solution: Optimized for VoLTE

The Cisco Policy solution is designed with VoLTE requirements in mind. The three-tier architecture shown in Figure 3 simplifies growth and reduces latency.

- Input/output layer (blue): Distributes policy requests to multiple virtualized applications. This layer includes a load-balancing feature, eliminating the need for a DRA.
- Application layer (green): Applies the policy management logic, requesting session information from the persistence layer.
- Persistence layer (red): Handles session information.

Figure 3. Three-Tier Architecture Simplifies Growth and Reduces Latency



All applications in the Cisco solution are virtualized. They do not require specialized servers. Instead, they operate on any standards-based server, with popular hypervisors, keeping costs down.

As your subscriber base grows, the policy management solution scales easily. All it takes is adding new application nodes. It is not necessary to change the persistence layer, modify the network, add a DRA, or rewrite your load-balancing software. A document-based database reduces latency while also avoiding licensing costs. Table 1 summarizes the benefits.

Table 1. Benefits of Cisco Policy Solution: Lower Costs, Better Service Quality, Easier Growth

Business Benefit	Reason
High Quality of Experience, Preventing Churn	Support hundreds of thousands of transactions per second, with average latency of less than 15 milliseconds.
Lower Capital and Operational Expense	Avoid cost of DRA server. Deploy one backup node for all policy nodes (n+1 redundancy) instead of one backup node for every policy node (1+1 redundancy). N+1 redundancy is possible because the Cisco solution uses stateless virtual machines in the application layer. Use commercial off-the-shelf (COTS) blade servers for virtualized applications. Eliminate database software licensing costs by using document-oriented database.
Ease of Adding Capacity	Add or subtract virtual servers as transaction volume increases or decreases. Take advantage of built-in load balancing. Scale linearly, with n+1 redundancy.
Resiliency	Deploy in geographically redundant architecture, with either an active-standby or active-active configuration.

Conclusion

VoLTE changes the game for policy management. Traditional, stateful policy management applications cannot meet the challenge. Their one- or two-tier application architectures make it difficult and costly to scale as transaction volume grows.

Cisco Policy solution lets you introduce VoLTE with confidence. Using separate applications for PCRF and session persistence makes it easier to scale. The solution reduces latency by eliminating the need for a DRA. And it helps you meet subscriber expectations for availability.

For More Information

To learn more about Cisco Policy solution visit: <http://www.cisco.com/c/en/us/products/wireless/policy-suite-mobile/index.html>.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)