

# Using Service Brokers for Accessing Backend Servers for Web Applications\*

Huamin Chen and Prasant Mohapatra  
Department of Computer Science  
University of California, Davis, CA 95616.  
Email: {chenhua, prasant}@cs.ucdavis.edu.

## Abstract

Web service infrastructures usually are comprised of front-end Web servers that accept requests and process them, and backend servers that manage data and services. Current Web servers use various API sets to access backend services. This model does not support service differentiation, overload control, caching of contents generated by backend servers. We have proposed a framework for using service brokers to facilitate these features. Service brokers are software agents that are the access points to backend services in Web servers. Unlike the current API-based scheme where accesses to backend services are through stateless and isolated APIs, in service broker framework, they are undertaken by passing messages to service brokers who gather all the requests and intelligently process them. We have prototyped this framework and validated its function in providing request clustering and service differentiation in accessing backend services. In addition, the performance in terms of the processing time is enhanced by this approach.

Keywords: HTTP, Web services, service broker, service differentiation, overload control, dynamic content caching.

## 1 Introduction

Web servers have established their presence and usage in a variety of environment. More and more servers are being deployed for complex service environments, which also involve a variety of auxiliary servers. The platform independence and universal accessibility of Web servers have been leveraged to access other services like database, mail, and directories. Web services like Microsoft .NET initiatives push such practice even further by facilitating more services accessibility through Web interfaces. It is conceivable that future Web servers will involve even more heterogeneous

auxiliary service providers (hereafter, referred to as backend servers) to serve various tasks. Most large Web servers include a set of front-end servers that receive the requests from the clients. The requests are served by accessing a set of backend servers, which provide database, directory services, secure transactions, and other services. A schematic diagram of a typical Web server environment is shown in Figure 1. It shows two applications running on a front-end Web server that has access to the backend servers.

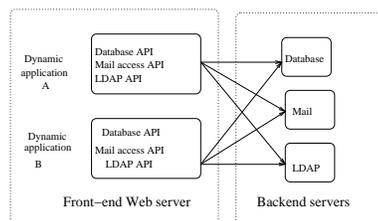


Figure 1. API paradigm.

Backend servers can be categorized as tightly coupled or loosely coupled based on their connectivity and ownership with the initiating Web servers. Tightly coupled servers, like database and directory servers, are closely connected, usually in the same LAN, to the Web servers and belong to the same administrative authority. Tightly coupled servers are usually reliable and of high capacity. Loosely coupled servers represent Web servers belonging to other owners, which are not under control of the request initiating front-end servers. Web syndicates like My.Yahoo! and My.Netscape belong to loosely coupled Web servers. In accessing their services, the requests and response traffic must traverse WAN networks, which may incur higher latency and jitters than LANs. In more security-sensitive applications, authentication must proceed before further transactions. Since the loosely coupled servers are shared resources, service guarantee becomes an outstanding problem. We envision that in the future such services would be contract-based such that the service availability is honored only when the incoming traffic are within the contracted

\*This research was supported in part by the National Science Foundation through the grants CCR-0296070 and ANI-0296034.

specifications. Loosely coupled services present a business model that has been existing in the current society. For instance, a travel agency has no sole control over airlines' ticketing services. Rather it contacts multiple airlines and selects the best deals for the customers.

The connectivity distinction between the two categories exposes different performance concerns. For tightly coupled services, the major performance issue is how to engineer the capacity to satisfy the front-end Web servers' requests. The in-equivalence of the computational complexity and different levels of replicability between the front-end Web and backend servers make backend servers the likely performance bottleneck. For instance, a search operation involves traversal of database tables with many comparison operations, which only results in a few lines of output that are rendered in HTML pages. Moreover, many companies use centralized backend servers to serve multiple front-end Web servers to reduce the investment on the expensive backend servers. Thus the priority in this environment is how to prevent and control overload at the backend servers. Similar concerns were expressed in [12]. In the loosely coupled environment, the volatile network conditions and the contract constraints demands efficient access schemes to ensure service availability and reliability.

In this paper we introduce a new concept and propose using *service brokers* to access backend services. Service brokers are software agents that act as access points to backend services in Web servers. Unlike the current API-based scheme where accesses to backend services are through stateless and isolated APIs, in the proposed scheme, backend services are accessed by passing messages to service brokers who gather all the requests and intelligently process them. In the tightly coupled environment, the service brokers can selectively drop some requests to reduce backend servers' load while facilitating service discrimination and QoS provision. In the loosely coupled environment, the brokers intelligently cluster the requests, cache the responses and prefetch the next possible queries in idle periods. We have prototyped this scheme and validated the feasibility in the testbed. We found that by properly clustering the backend server accesses, service brokers can significantly reduce the response time. In the service differentiation experiments, we demonstrate notable scalability improvement through fidelity variations. The experimental results depict a significant performance benefit through the usage of service brokers for backend services.

The rest of the paper is organized in this way. Section 2 discusses the current API-based framework and its drawbacks. Section 3 introduces the service broker model and its advantages in various aspects. Section 4 presents the an implementation of the service broker model and its ability in request clustering and service differentiation. It is observed that by varying response fidelity in different QoS levels,

service brokers can improve responsiveness and scalability. The related works were presented in Section 5 followed by the concluding remarks in Section 6.

## 2 API-Based Backend Service Access

Currently dynamic Web applications are in forms of CGI executables that run in separate processes or use scripting languages like PHP and ASP which usually run in Web server processes. They access backend servers (Database, LDAP, mail or even other Web servers) through specific APIs like socket, ODBC or modules like COM. The APIs reside inside the application process space and share no information among different processes. The paradigm adopted in the contemporary Web servers is illustrated in Figure 1. In this paradigm, applications A and B need to access database, mail and directory servers. They use the respective API sets to accomplish the tasks: in order to access database server, a connection needs to be established before any queries are initiated which are followed by the connection tear-down; similar procedures are applied in accessing the directory server. Since A and B are in different process space, their use of API sets is independent and they do not share anything. Thus even though A and B may access database or directory service simultaneously, they each need to establish a connection before using any services.

The drawbacks of this paradigm are:

- QoS is not guaranteed: Accesses to backend service are served on the FCFS basis. Unless QoS specifications propagate through all backend servers, there is no guarantee that they can be honored. Currently most of the backend servers like database servers, directory servers, and file servers do not provide QoS support. It is also likely that heterogeneous backend servers may have different QoS notations. For example, the file servers may cluster requests whose accesses are in adjacent disk layout. Database servers may cluster queries that access the same table. Thus this architecture may not be able to enforce QoS specifications throughout the entire systems and is also subject to the priority inversion problem<sup>1</sup>.
- Hot spot unawareness: Although overload control in Web servers has received increasing attention recently, there has been limited research work on backend servers' overload control. In fact, backend servers are more likely to be overloaded due to its less replicability (it is usually very expensive to maintain replicated backend servers.). When the traffic to the same backend server is beyond its capacity, a hot spot is

---

<sup>1</sup>A low priority request could get service earlier than a high priority request.

generated and this backend sever is likely to become bottleneck of the entire request handling process. Hot spots generated in backend servers are at most known to those who are using the service. Other processes are unaware of the overload due to the independence of the request handling processes. Thus overload could spread to other processes. The increased number of processes that are trapped in hot spots could impose serious threat to the overall server performance in some server architectures. For example, in Apache Web server, each request is handled by a dedicated server process. In order to process incoming requests, more child processes must be forked if others are busy. Thus processes trapped in accessing overloaded backend resources essentially exacerbate the overall performance.

- Accesses are isolated, and thus not optimized globally: Most of API libraries for backend server accesses do not share states or results among individual instances. Each application send requests and launch I/O operations separately even for identical operations. For instance, a database access involves connection establishment and tear-down, query command transmission, and result retrieval. Two API-based applications need to traverse these steps individually. But the connection establishment and tear-down could be shared.

### 3 Service Broker Paradigm

Motivated by the shortcomings of the contemporary model of backend accesses, we propose a service broker model in which instead of using APIs to access backend servers, Web servers pass requests onto intermediate dedicated processes (referred as service brokers). The schematic architecture is illustrated in Figure 2. In this architecture, dynamic applications A and B do to have to invoke APIs to access backend services. They only pass messages to individual service brokers in some formats that contain their QoS specification and queries. Service brokers receive, sort and rewrite these messages according to their QoS levels and carry out the real query across connections. These connections can be established in advance to reduce the setup overhead, which is especially beneficial to loosely coupled environment. Upon receiving replies from backend servers, service brokers send the results to the dynamic applications and make a local copy, if possible, to serve similar requests. When hot spot occurs, the brokers can take appropriate actions globally across all the requests. In short, service brokers depart from current decentralized backend server access to a moderated control model.

As illustrated in Figure 3, service brokers are independent of the Web application logic and are built on top of the API sets. It is per service based. Though schematically

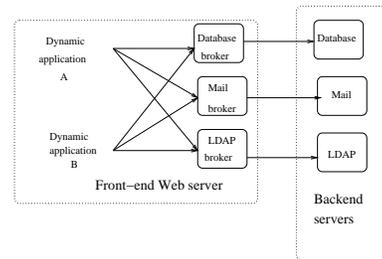


Figure 2. Service broker framework.

similar to COM and Enterprise Java Bean, service brokers have more meaningful attributes.

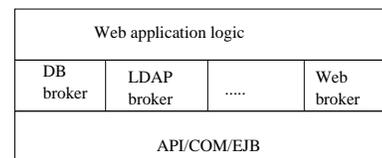


Figure 3. Service broker Layer

At the first glance, this new proposal appears to incur extra overhead for inter-process communication between service brokers and web processes. However, to access backend services, such overhead is insignificant. For a database access, database connection and tear-down, which are required in API model for each access, would be more expensive than inter-process communication. In the proposed approach, DB brokers maintain persistent connection thus saving the cost of connection setup.

The service broker framework has the following advantages:

- Accesses can be clustered and optimized: Multiple query clustering and optimization [10] has been studied in database systems. Service brokers can provide similar optimization among requests in absence of the backend server support. For instance, two separate accesses to a remote Web server to get page 1.html and 2.html can be combined using MGET command [7] as “MGET URI:1.html URI:2.html” at the broker and the results are appropriately split and sent to the request initiators. It appears that channel multiplexing may increase the workload at service brokers, which usually reside in the front-end Web servers. However, the transfer of computing load from backend servers to the front-end Web servers is viable since the front-end Web servers are easily replicable. Moreover, the same connection to backend servers can be multiplexed. In contrast to using independent connections to access backend servers in the API model, a single connection between the service broker and the backend server

can be multiplexed to serve multiple applications and thus reduce the connection overhead while processing queries in bulk.

- **QoS awareness:** Service brokers are able to consistently honor the QoS priorities without propagating QoS specification to backend servers. QoS rules are fed into all the service brokers. Based on these rules, service brokers reshuffle the queued requests and schedule according to their priorities. When traffic intensity of QoS classes exceed their limits, their requests are dropped and other classes are not affected. Therefore lower priority requests give way to higher priority classes, thus avoiding the priority inversion problem. We demonstrate this functionality in Section 4.
- **Backend server overload control:** Unlike API-based architecture where each backend service access is unrelated, service brokers process all the requests and thus are aware of the states of the associated backend servers. Service brokers can notify request schedulers about the onset of hot spots or respond to the requests with lower fidelity results which would indicate that the system is busy. In some cases, it enables the use of cached results from the previous queries.
- **Caching of query results:** Since service brokers receive all the query results from the same backend servers, they can cache some of the results to serve similar requests. For example, consider an online Web site that provides movie schedules. All the schedule information is stored in a database. In the peak time, there would be a lots of requests for the same movie schedule. If the results are not cached, the database has to process the same query repeatedly and will contribute to the response delay. In contrast, service brokers can be configured to cache the popular query results and promptly return them to the front-end Web servers. This approach reduces the number of requests to the backend servers and reduces the response time. This feature is especially beneficial in loosely coupled environment where accesses to backend service need to traverse high latency networks.
- **Prefetching:** Service brokers enable forecasting of the next possible queries and prefetching the necessary information. For instance, a news provider website periodically updates the online headlines. Service brokers can be synchronized to prefetch them when the server load is not high. So the requests for the news can be served immediately without accessing the backend servers.
- **Multitasking:** Requests that consist of independent heterogeneous tasks can send simultaneous messages

to service brokers which run in parallel to speed up processing. For example, a web syndicate like My.Yahoo composes contents from different and independent providers. Thus the page generator can send requests in parallel to service brokers that are associated with individual providers. The content retrievals can be overlapped to reduce the overall response time.

- **Load balancing:** Load balancing have been widely used to improve front-end Web servers' performance [6, 5]. The basic idea is to select a candidate server and distribute the workload across multiple servers. Similar idea is also applicable to backend service accesses. In the API-based architecture, since no state information is shared in individual accesses, it can only work in a speculative manner. The service brokers can track the traffic and monitor their workload and accurately distribute the workload among the backend servers to achieve a balanced load.
- **Amortized context switching:** Accesses to backend servers are done in bulk at service brokers to reduce the number of context switchings. The service brokers can reside separately from the Web servers to facilitate the overall system optimization.

## 4 Experimental Evaluation

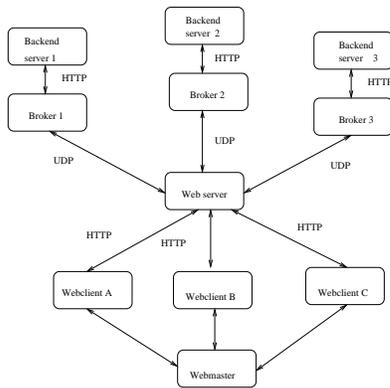
We have prototyped the service broker framework to demonstrate its feasibility and ability in service differentiation. The testbed configurations and results are discussed in this section.

### 4.1 Testbed Configuration

The prototype of the service broker scheme in the *boa* web server [4]. *Boa* is an open sourced, light-weighted, and efficient Web server. Similar implementation can be done on other Web servers. The objective of this experiment is to demonstrate the service broker's capabilities for service differentiation and the performance benefits.

As illustrated in Figure 4, in the testbed, there were 3 service brokers (brokers 1, 2 and 3), each connected to a Web server as backend server (backend server 1, 2 and 3, all ran Apache Web servers.). The brokers and the front-end Web server exchange request and response messages through lightweight UDP. The brokers communicate with the backend servers in HTTP. All the front and backend servers and brokers ran on Redhat Linux 7.2. The client side used WebStone 2.5 [11] on 4 Sun UltraSparc workstations. Three workstations A, B, and C were designated as Web clients with QoS levels 1, 2 and 3, respectively.

The backend services provided by each backend servers are CGI requests with bounded processing time. The processing time of each of the services is 1, 2 and 3 seconds at the backend servers 1, 2 and 3, respectively. The QoS specification used in the testbed is just a binary mode of forward or drop: QoS level  $i$  means that the request is forwarded to the backend servers if the number of the outstanding requests is  $(10 - x) * 10\%$  of the threshold. For instance, a request with QoS level 2 is allowed to send query to the corresponding backend server only when the number of outstanding requests is below 80% of the threshold. The thresholds at each broker were set to be 20, i.e. at most 20 requests are allowed to be outstanding in each of the backend servers. The thresholds of queues ensure bounded queueing time. The maximum number of server processes in each of the backend end Web servers is set to be 5, therefore only 5 requests can be processed simultaneously in each of these servers and the rests are queued. If a request was dropped at the broker, a short message was sent to the front-end Web server immediately. Therefore, longer the processing time a request undergoes, higher the fidelity it receives.



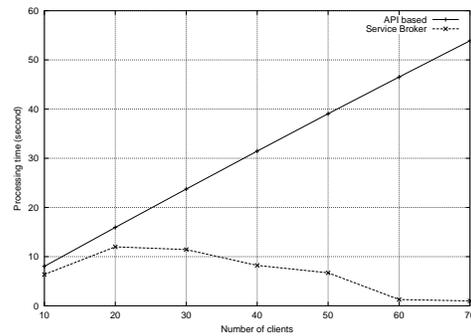
**Figure 4. Service differentiation testbed configuration.**

## 4.2 Experimental Results

We set up an experiment to demonstrate the QoS assurance provided by service brokers. The backend servers were assured to be QoS-unaware. A normal Web request consists of 3 stages which take approximately 6 seconds to complete. In the service broker scheme, if a request's QoS level does not meet the current load status, a low fidelity response is replied immediately. Thus the longer is the processing time of a request, the higher is the fidelity of service it receives.

The experiments were conducted for both API-based accesses and service broker-based configurations. Figure 5

presents the results reported by WebStone. It is observed that in API-based accesses, the average processing time is linear with respect to the number of requests. In the service broker model, the processing time first rises when the number of clients is small so service brokers can admit most of the requests to access backend servers. When the number of clients increases, more and more requests in lower priority classes are dropped at the service brokers and they are informed promptly without any backend service, so the processing time declines. Lower priority requests are dropped to give way to process requests in higher priority classes. Thus the fidelity of each QoS class is differentiated. In this regard, the priority inversion abnormality is avoided.

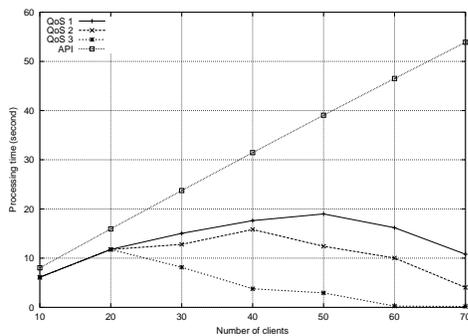


**Figure 5. Processing time of API and service broker based settings.**

Figure 6 depicts the average processing time of each QoS class under various load. The requests with higher QoS level experienced longer processing time, which means that the fidelity of the response is higher. For each curve, an increase in processing time is observed followed by a decline when the number of clients reaches a certain point. The reasons behind this phenomena are: when the number of clients was low (from 10 to 20), only a few requests were queued in service brokers, thus the increase of requests could be accommodated; after some points, the traffic intensity exceeded the capacity and some lower prioritized requests were responded by service brokers with lower fidelity messages without being forwarded to the backend servers. High priority requests still enjoy high fidelity service but the increase in traffic intensity contributes to the queueing time. When the number of clients exceeded 50, more requests from the high QoS classes were dropped at the service brokers and thus their average processing time declined as well.

## 5 Related Works

The related work can be categorized into component-based systems and Web server clustering.



**Figure 6. Average processing time for each QoS level.**

Pupeteer [8] is a component-based system for mobile devices. It shares the concepts of service brokers. In Pupeteer, network activity of mobile applications is governed by a common middleware which determines bandwidth allocation, content caching, and content transcoding. While the service broker framework applies to different system environment: interaction between front-end Web and backend servers. In addition to the aspects addressed in Pupeteer, more issues that are particular to Web services can be solved in this framework.

Most of the studies on multiple servers environment in Web research are focused on how to improve performance through load balancing techniques. SWEB [1] investigated how to use DNS to implement load balancing and corresponding support in operating systems. Cluster Reserves [3] proposed a technique to provide resource isolation among different classes. Locality-aware request dispatching algorithms were proposed in [9, 12] to enable requests with related URL to be sent to the same server to achieve locality. Our work is complementary to these front-end Web server research.

## 6 Conclusions

Web servers have been used widely for applications that need to access backend services. The current API-based access schemes are not well suited for various reasons. We proposed a service broker framework that facilitate performance improvement and QoS provisioning. In the service broker framework, every backend service access is carried out by middleware agents that reside between front-end Web servers and backend servers. Service brokers can cluster and optimize queries and cache the results according to a set of pre-defined rules. We have prototyped the framework to demonstrate their feasibility and evaluate their capacity to provide service differentiation. Experimental results demonstrated that service brokers can effectively use

QoS specifications and provide scalable performance and service differentiation.

## References

- [1] D. Andresen, T. Yang, V. Holmedahl, and O.H. Ibarra, "SWEB: Toward a Scalable World Wide Web Server on Multicomputers," In Proceedings of the 10th International Parallel Processing Symposium, Honolulu, Hawaii, USA , April 1996.
- [2] Apache HTTP Server Project, <http://www.apache.org>
- [3] M. Aron, P. Druschel, and W. Zwaenepoel, "Cluster Reserves: A Mechanism for Resource Management in Cluster-Based Network Servers," In Proceedings of the ACM SIGMETRICS 2000 Conference, Santa Clara, California , June 2000. ACM.
- [4] Boa Web server, <http://www.boa.org>.
- [5] H. Bryhni, E. Klovning, and O. Kure, "A Comparison of Load Balancing Techniques for Scalable Web Servers," IEEE Network, pages 58–64, July/August 2000.
- [6] V. Cardellini, M. Colajanni, and P. S. Yu, "Load Balancing on Web-server Systems," IEEE Internet Computing, 3(3):28–39, May/June 1999.
- [7] J. Franks, "An MGET proposal for HTTP," October 1994. WWW-talk mailing list. <http://www.webhistory.org/www.lists/www-talk.1994q4/0479.html>
- [8] E. D. Lara, D. Wallach, and W. Zwaenepoel, "Pupeteer: Component-based Adaptation for Mobile Computing," in Proceedings of the 3rd Usenix Symposium on Internet Technologies and Systems, March 2001.
- [9] V. S. Pai, M. Aron, G. Banga, M. Svendsen, P. Druschel, W. Zwaenepoel, and E. Nahum, "Locality-Aware Content Distribution in Cluster-Based Network Servers," Proc. of the Conf. on Architectural Support for Programming Languages and Operating Systems, San Jose, CA Oct. 1998.
- [10] T. Sellis, "Multiple Query Optimization", ACM Transactions on Database Systems, vol.13, no.1, pp.23-52, 1988.
- [11] WebStone, <http://www.mindcraft.com/webstone/>
- [12] X. Zhang, M. Barrientos, J. Chen, and M. Seltzer, "HACC: An Architecture for Cluster-Based Web Servers," in Proceedings of the Third USENIX Windows NT Symposium. July 1999.