

---

## A hybrid connector for efficient web servers

---

David Carrera\*, Vicenç Beltran, Jordi Torres  
and Eduard Ayguadé

Barcelona Supercomputing Center (BSC),  
Technical University of Catalonia (UPC), Barcelona, Spain  
E-mail: dcarrera@ac.upc.edu  
E-mail: vbeltran@ac.upc.edu  
E-mail: torres@ac.upc.edu  
E-mail: eduard@ac.upc.edu

\*Corresponding author

**Abstract:** In this paper we introduce a novel web server architecture that combines the best aspects of both the multithreaded and the event-driven architectures, the two major existing alternatives, to create a server model that offers an improved performance in terms of user session completions without losing the natural ease of programming characteristic of the multithreading paradigm. We describe the implementation of this architecture on the Tomcat 5.5 server and evaluate its performance. The obtained results demonstrate the feasibility of the hybrid architecture and the performance benefits that this model introduces for e-commerce applications.

**Keywords:** web servers; performance evaluation; java; network-based computing; scalability.

**Reference** to this paper should be made as follows: Carrera, D., Beltran, V., Torres, J. and Ayguadé, E. (2008) 'A hybrid connector for efficient web servers', *Int. J. High Performance Computing and Networking*, Vol. 5, Nos. 5/6, pp.323–330.

**Biographical notes:** David Carrera received the Engineering Degree in Computer Science in 2002 and the PhD in Computer Science in 2008, both from the Technical University of Catalonia (UPC), where he is an Assistant Professor. He is also an associate researcher at the Barcelona Supercomputing Center (BSC). His research interests are focused on the integrated performance management of virtualised data centers that host applications with differentiated QoS criteria.

Vicenç Beltran received the Engineering Degree in Computer Science in 2004 and the PhD in Computer Science in 2008, both from the Technical University of Catalonia (UPC). He is currently a senior researcher at the Barcelona Supercomputing Center (BSC). His research interests cover the management of virtualised environments and the exploitation of emerging hardware.

Jordi Torres received the Engineering Degree in Computer Science in 1988 and the PhD in Computer Science in 1993, both from the Technical University of Catalonia (UPC), where he is an Associate Professor. He is the manager of the Autonomic Systems and e-Business Platforms at the Barcelona Supercomputing Center (BSC). His research interests include applications for parallel and distributed systems, web applications and performance analysis and prediction tools.

Eduard Ayguadé received the Engineering Degree in Telecommunications in 1986 and the PhD in Computer Science in 1989, both from the Technical University of Catalonia (UPC) where he is a Professor since 1997. He is the Associate Director of the Computer Science Department at the Barcelona Supercomputing Center (BSC). His research interests cover the areas of processor micro-architecture and ILP exploitation and parallelising compilers for highperformance computing systems.

---

### 1 Introduction

Most e-commerce applications are distributed applications based on the well-known client/server paradigm that reside mostly in an application server and that are usually accessed by a remote thin web client. The communication

protocol between the server and the client is the Hypertext Transfer Protocol (HTTP), which in the server side is parsed and processed by a component of the application server that is commonly known as the web container. Most generally, a web container can be considered as a web server that can support some language extensions to create more flexible web applications.

The performance of an e-commerce application can be measured in technical terms or following business indicators (see Keller and Ludwig, 2002, for further discussion on this topic). Usually, the technical measurement units for a web server are replies per second (throughput) and the response time. In contrast, most performance metrics based on business indicators can be reduced to one: profit. The revenues a website can generate are directly related to the amount of commercial transactions the website can complete. In general, one commercial transaction is completed when user browsing session successfully finishes. Therefore, the web container architectures must be designed to obtain a high performance in user sessions units in order to be on the side of the interests of e-commerce applications.

The architectural design of most currently existing web containers follow the multithreading paradigm (Apache and Tomcat Jakarta Project are widely extended examples) because it leads to a natural ease of programming and simplifies the development of the web container logic. Unfortunately, this model is not especially indicated to obtain a high performance in terms of user session completions. Alternatively, an event-driven model (used in Flash (Pai et al., 1999) and in the SEDA (Welsh et al., 2001) architecture) can be adopted to develop a web server container. This model solves some of the problems present in the multithreaded architecture but transforms the development of the web container into a hard task.

In this paper we introduce a new hybrid web server architecture that exploits the best of each one of the discussed server architectures. With this hybrid architecture, an event-driven model is applied to receive the incoming client requests. When a request is received, it is serviced following a multithreaded programming model, with the resulting simplification of the web container development associated to the multithreading paradigm. When the request processing is completed, the event-driven model is applied again to wait for the client to issue new requests. With this, the best of each model is combined and, as it is discussed in following sections, the performance of the server is remarkably increased in terms of user session completions.

The rest of the paper is structured as follows: Section 3 discusses the characteristics of the two architectures involved in the creation of a hybrid web server, Section 4 describes the implementation details of the hybrid web server evaluated on this paper, later, in Section 5, we present the execution environment where the experimental results presented in this work were obtained. Finally, Section 6 presents the experimental results obtained in the evaluation of the hybrid web server and Section 7 gives some concluding remarks and discusses some of the future work lines derived from this work.

## 2 Related work

The use of event-driven architectures, instead of the traditional multithreading approach, for web servers is an already explored area. Flash (Pai et al., 1999) is

an asymmetric multi-process event-driven web server which exploits the creation of a set of helper processes to avoid thread-blocking in the main servicing process. Haboob (Welsh et al., 2001) is also an event-driven web server, based on the concepts of staged event-driven servers introduced by SEDA (Welsh et al., 2001). JAWS (Hu and Schmidt, 1999) web server uses the Proactor (Hu et al., 1998) pattern to easily construct an event-driven concurrency mechanism. The architecture proposed and evaluated in this paper evaluates an hybrid architecture, which combines the benefits of both the event-driven and the multithread architectures. None of the previous works evaluate mixed solutions as a next step in the performance of web servers.

The hybrid architecture is completely compatible with any other performance optimisations explored for the web servers, such as request scheduling and resource reservation and provisioning, since the architectural redefinition presented in this paper introduces a new degree in easiness of web container development and a better network I/O resource management, but does not modify the way how the other system resources (CPU, disk . . .) are exploited. In summary, the benefits introduced by this architecture can be inherited by any additional performance improvement that can be introduced in a web server container.

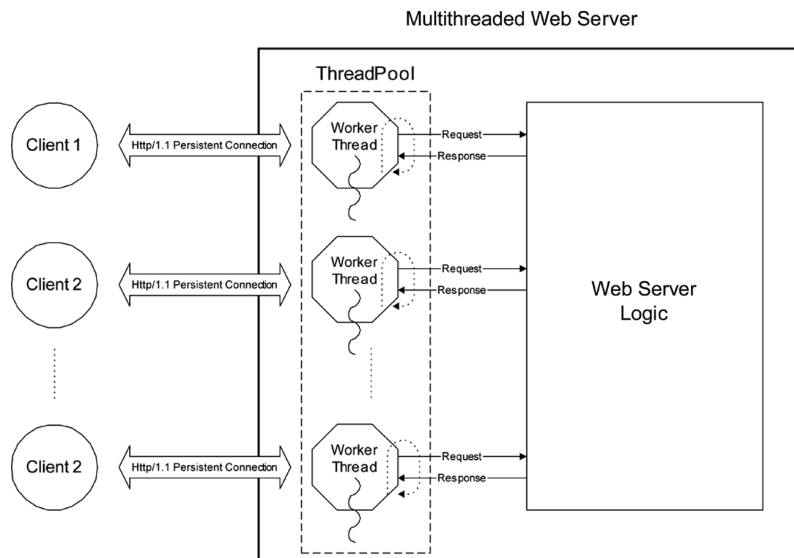
## 3 Web server architectures

There are multiple architectural options for a web server design, depending on the concurrency programming model chosen for the implementation. The two major alternatives are the multithreaded model and the event-driven model. In both models, the work tasks to be performed by the server are divided into work assignments that are assumed each one by a thread (a worker thread). If a multithreaded model is chosen, the unit of work that can be assigned to a worker thread is a client connection, which is achieved by creating a virtual association between the thread and the client connection that is not broken until the connection is closed. Alternatively, in an event driven model the work assignment unit is a client request, so there is no real association between a server thread and a client.

### 3.1 Multithreaded architecture with blocking I/O

The multithreaded programming model leads to a very easy and natural way of programming a web server. The association of each thread with a client connection results in a comprehensive thread lifecycle, started with the arrival of a client connection request and finished with the connection close. This model is especially appropriate for short-lived client connections and with low inactivity periods, which is the scenario created by the use of non persistent HTTP/1.0 connections. A pure multithreaded web server architecture is generally composed by an acceptor thread and a pool of worker threads. The acceptor thread is in charge of accepting new incoming connections, after what each established connection is assigned to one

Figure 1 Multithread architecture diagram



thread of the workers pool, which will be responsible of processing all the requests issued by the corresponding web client. A brief operation diagram for this architecture can be seen on Figure 1.

The introduction of connection persistence in the HTTP protocol, already in the 1.0 version of the protocol but mainly with the arrival of HTTP/1.1, resulted in a dramatic performance impact for the existing multithreaded web servers. Persistent connections, which means connections that are kept alive by the client between two successive HTTP requests that in turn can be separated in time by several seconds of inactivity (think times), cause that many server threads can be retained by clients even when no requests are being issued and the thread keeps in idle state. The use of blocking I/O operations on the sockets is the cause of this performance degradation scenario. The situation can be solved increasing the number of threads available (which in turn results in a contention increase in the shared resources of the server that require exclusive access) or introducing an inactivity timeout for the established connections, that can be reduced as the server load is increased. When a server is put under a severe load, the effect of applying a shortened inactivity timeout to the clients lead to a virtual conversion of the HTTP/1.1 protocol into the older HTTP/1.0, with the consequent loss of the performance effects of the connection persistence.

In this model, the effect of closing client connections to free worker threads reduces the probability for a client to complete a session to nearly zero. It is especially important when the server is under overload conditions, where the inactivity timeout is dynamically decreased to the minimum possible in order to free worker threads as quickly as possible, which provokes that all the established connections are closed during think times. This causes a higher competition among clients trying to establish a connection with the server. If we extend it to the length of a user session, we obtain that the probability of finishing it successfully under this architecture is still much lower than the probability of establishing each one of the connections it is composed of, driving the server to obtain

a really low performance in terms of session completions. This situation can be alleviated increasing the number of worker threads available in the server, but this measure also produces an important increase in the internal web container contention with the corresponding performance slowdown.

### 3.2 Event-driven architecture with non-blocking I/O

On the other hand, the event-driven architecture completely eliminates the use of blocking I/O operations for the worker threads, reducing their idle times to the minimum because no I/O operations are performed for a socket if no data is already available on it to be read. With this model, maintaining a big amount of clients connected to the server does not represent a problem because one thread will never be blocked waiting a client request. With this, the model detaches threads from client connections, and only associates threads to client requests, considering them as an independent work units. An example of web server based on this model is described in Pai et al. (1999), and a general evaluation of the architecture can be found in Welsh et al. (2001).

In an event driven architecture, one thread is in charge of accepting new incoming connections. When the connection is accepted, the corresponding socket channel is registered in a channel selector where another thread (the request dispatcher) will wait for socket activity. Worker threads are only awakened when a client request is already available in the socket. When the request is completely processed and the reply has been successfully issued, the worker thread registers again the socket channel in the selector and gets free to be assigned to new received client requests. This operation model avoids worker threads to keep blocked in socket read operations during client think times and eliminates the need of introducing connection inactivity timeouts and their associated problems.

A remarkable characteristic of the event-driven architectures is that the number of active clients connected to the server is unbounded, so an admission control (Chen and Mohapatra, 2002) policy must be implemented. Additionally, as the number of worker threads can be very low (one should be enough) the contention inside the web container can be reduce to the minimum.

### 3.3 Hybrid architecture

In this paper we propose a hybrid architecture that can take benefit of the strong points of both discussed architectures, the multithreaded and the event driven. In this hybrid architecture, the operation model of the event-driven architecture is used for the assignment of client requests to worker threads (instead of client connections) and the multithreaded model is used for the processing and service of client requests, where the worker threads will perform blocking I/O operations when required. This architecture can be used to decouple the management of active connections from the request processing and servicing activity of the worker threads. With this, the web container logic can be implemented following the multithreaded natural programming model and the management of connections can be done with the highest possible performance, without blocking I/O operations and reaching a maximum overlapping of the client think times with the processing of requests. A brief operation diagram for this architecture can be seen on Figure 2.

In this architecture the acceptor thread role is maintained as well as the request dispatcher role from the pure event-driven model and the worker thread pool (performing blocking I/O operations when necessary) from the pure multithreaded design. This makes possible for the hybrid model to avoid the need of closing connections to free worker threads without renouncing to the multithreading paradigm. In consequence, the hybrid architecture makes a better use of the characteristics introduced to the HTTP protocol in the 1.1 version,

such as connection persistence, with the corresponding reduction in the number of client reconnections (and the corresponding bandwidth save). Additionally, and as it has been discussed for the event-driven architectures, some kind of admission control policy must be implemented in the server in order to maintain the system in an acceptable load level.


## 4 Hybrid architecture implementation

To validate the proposed hybrid architecture we have implemented it inside of Tomcat 5.5, a widely extended web container. Tomcat 5 is built in the top of the Java platform, which provides non blocking I/O facilities across different operating systems in its Non Blocking I/O (NIO) API. The implementation has been done modifying a pluggable component of the server generally named connector which is the server component in charge of handling communications with the client

Tomcat is an open-source servlet container developed under the Apache license. Its primary goal is to serve as a reference implementation of the Sun Servlet and JSP specifications, and to be a quality production servlet container. For the scope of this paper, two major components of Tomcat are considered: Coyote and Catalina. Coyote is the default Tomcat connector and deals with client connection request parsing and thread pooling. Catalina is a servlet container, and implements most of the web container logic. The implementation of the hybrid server architecture in Tomcat only affects Coyote.

Coyote follows a connection service schema where, at a given time, one thread (an `HttpProcessor`) is responsible of accepting a new incoming connection on the server listening port and assigning to it a socket structure. From this point, this `HttpProcessor` will be responsible of attending and parsing the received requests through the persistent connection established with the client, while another `HttpProcessor` will continue accepting new connections. `HttpProcessors` are commonly chosen

Figure 2 Hybrid architecture diagram



from a pool of threads in order to avoid thread creation overheads. A connection timeout is programmed to close the connection if no more requests are received in a period of time. When a request is parsed, Coyote requires Catalina to process the request and to send the corresponding response to the client.

The implementation of the hybrid architecture changes the original Coyote threading and I/O model. One thread is in charge of accepting and registering, through a NIO selector, new incoming connections. When a registered connection becomes active (i.e., it has data available to read so a read operation over the socket will not be blocking), it is dispatched by the selector thread to small pool of `HttpProcessor` threads. Each `HttpProcessor` services only one request for each assigned active connection. The request is read and parsed, always without blocking the thread, and it is sent to Catalina who processes the request. When the request is finished the connection is re-registered on the selector and the thread is sent back to the pool until a new active connection is assigned to it.

This implementation presents a major drawback when the system becomes overloaded because in this situation the acceptor thread allows new connections to enter the server faster than Catalina can service them, what results in a quick growth of the number of concurrent connections, which in turn causes a severe response time degradation. In consequence the number of client timeouts grows and the throughput decreases. To avoid this problem we have introduced a simple but effective admission control mechanism (similar to the backpressure technique described in Welsh et al. (2001)), that prevents the acceptor thread from accepting new connections while all `HttpProcessors` are busy.

## 5 Testing environment

The hardware platform for the experiments presented in this paper is composed of a 4-way Intel Xeon 1.4GHz with 2GB RAM to run the web servers and a 2-way Intel XEON 2.4GHz with 2GB RAM to run the benchmark clients. For the benchmark applications that require the use of a database server, a 2-way Intel XEON 2.4GHz with 2GB RAM was used to run MySQL v4.0.18, with the MM.MySQL v3.0.8 JDBC driver. All the machines were running a Linux 2.6 kernel, and were connected through a switched Gbit network. The SDK 1.5 from Sun was used to develop and run the web servers.

The servers were tested in two different scenarios, one to evaluate the server performance for a static content application and another for a dynamic content environment. The workload for the experiments was generated using a workload generator and web performance measurement tool named `Httpperf` (Mosberger and Jin, 1998). This tool allows the creation of a continuous flow of HTTP requests issued from one or more client machines and processed by one server machine: the System Under Test (SUT). The configuration parameters of the benchmarking tool used

for the experiments presented in this paper were set to create a realistic workload, with non-uniform reply sizes, and to sustain a continuous load on the server. One of the parameters of the tool represents the number of clients/s, i.e., the load. Each emulated client opens a session with the server. The session remains alive for a period of time, called session time, at the end of which the connection is closed. Each session is a persistent HTTP connection with the server, used by the client to repeatedly send requests, some of them pipelined. The requests issued by `httpperf` were extracted from the surge (Barford and Crovella, 1998) workload generator for the static content scenario and from the RUBiS (Amza et al., 2002) application for the dynamic content environment.

A static content application is characterised by the short length of the user sessions as well as by the low computational cost of each request to be serviced. This is the scenario reproduced with the Surge (Barford and Crovella, 1998) workload generator for these experiments. The request distribution produced by Surge is based on the observation of some real web server logs, from where it was extracted a data distribution model of the observed workload. This fact guarantees that the used workload for the experiments follows a realistic load distribution.

A dynamic content application is characterised by the long length of the user sessions (an average of 300 requests per session in front of the six requests per session for the static workload) as well as by the high computational cost of each request to be serviced (including embedded requests to external servers, such as databases). For our experiments, the chosen dynamic content benchmark application was (Amza et al., 2002) Rice University Bidding System (RUBiS), in its 1.4 servlets version. RUBiS implements the core functionality of an auction site: selling, browsing and bidding. The workload distribution generated by `Httpperf` was extracted from the RUBiS client emulator, which uses a Markov model to determine which subsequent link from the response to follow. Each emulated client waits for an amount of time, before initiating the next interaction, emulating on this way the “thinking times observed in real clients. `Httpperf` allows also configuring a client timeout. If this timeout is elapsed and no reply has been received from the server, the current session is discarded.

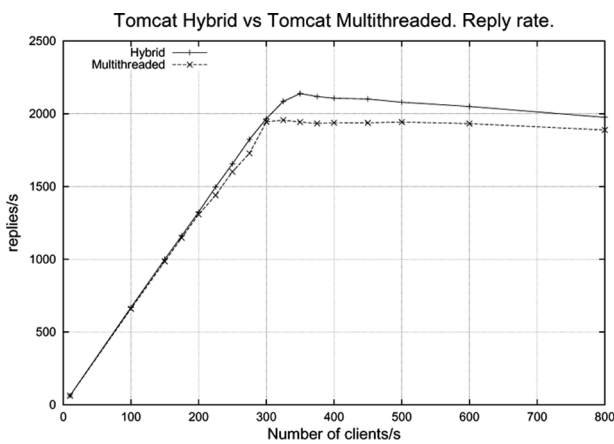
## 6 Experimental results

In our experiments we evaluate which are the performance characteristics for each web server architecture, the original multithreaded Tomcat and our hybrid Tomcat implementation, for the two different scenarios already introduced in Section 5: a static content application (Surge) and a dynamic content one (RUBiS). For the experiments, we configured `Httpperf` setting the client timeout value to 10 s, and used the configuration described in Section 5. Each individual benchmark execution had a fixed duration of 30 min for the dynamic content tests and 10 min for the static content experiments.

### 6.1 Static content

The first performance metric evaluated for this scenario is the obtained throughput for each architectural design, measured in replies per second. The results for both architectures are shown in Figure 3. In the figures, it can be seen that the multithreaded architecture obtains a slightly lower performance than the hybrid one in terms of throughput, but the results are so close that they can be considered equivalent. It is remarkable that the same throughput is obtained by the hybrid architecture with a thread pool size of only ten threads, while in the multithreaded architecture requires 500 threads to obtain the same result.

**Figure 3** Throughput comparison under an static content workload

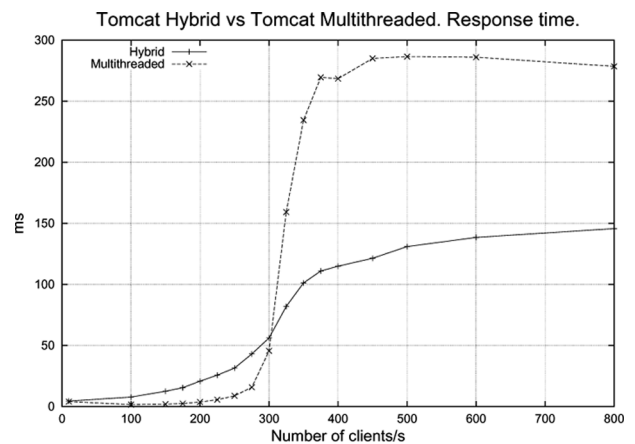


If we move from the throughput to the response time observed for each implementation, we can see that the hybrid architecture offers a clearly better result than the multithreaded one, as it can be seen in Figure 4. When the system is not saturated (under a load equivalent to 300 new clients per second), the response time for the multithreaded server is slightly better than for the hybrid design, possibly because of the overhead introduced by the extra operations that the hybrid server must do to register the sockets in the selector and to switch them between blocking and non-blocking mode when moving from multithreaded to the event-driven and reverse. This effect would be dispelled in a WAN environment, where the latencies are much higher than in the Gbit LAN used for the experiments. When the system is saturated, beyond 300 new clients/s, the benefits of the hybrid architecture turn up and the response time reduction with respect to the multithreaded version of the server is of about a 50%, moving from a 300 ms average response time for the multithreaded architecture to a 150 ms response time for the hybrid design.

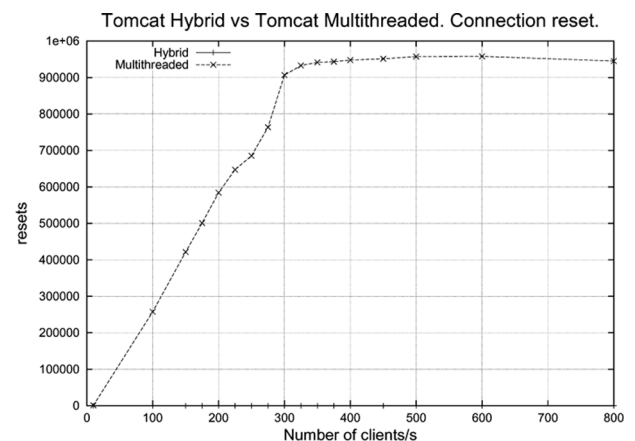
Another interesting difference between the behaviour of the two studied architectures can be observed in Figure 5. The figure shows the amount of connections that have been closed by the server because of a too long client inactivity period, causing the server timeout to expire and obligating the client to be reconnected to resume its session.

As it can be seen, while the hybrid architecture is not producing this kind of situation (zero errors are detected by the benchmark client), a high number of errors are detected for the multithreaded architecture. This situation can be explained by the need of the multithreaded design to free worker threads to make them available for new incoming connections. This causes that the client inactivity periods must be avoided by closing the connection and requiring the client to resume its session with a new connection establishment when necessary. In the hybrid architecture, that assigns client requests as work units to the worker threads instead of client connections, this situation is naturally avoided and the cost of keeping a client connection event in periods of inactivity is equivalent to the cost of keeping the connection socket opened. The effect of this characteristic for the hybrid architecture is that all the reconnections are eliminated.

**Figure 4** Response time under an static content workload



**Figure 5** Number of connections closed by the server by a timeout expiration



### 6.2 Dynamic content

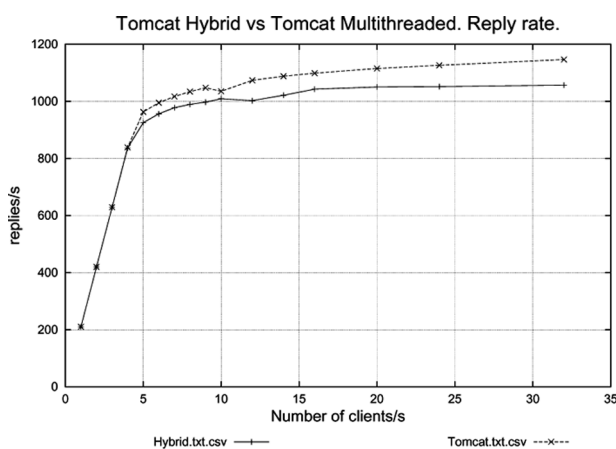
Dynamic content applications implement a higher complexity and more developed semantics than static ones, which is usually translated into longer user sessions and involves that the common performance metrics are partially redefined in terms of business concepts.

This means that e-commerce applications are more concerned about sales or business transactions than about more technical metrics such as the throughput or the response time offered by the server.

For this experiment, we consider that one of the most important metrics for an auction website (the scenario reproduced by RUBiS, see Section 5 for more details) is the number of user sessions that are completed successfully. Each user that can complete its navigation session represents a potential bid for an item and in consequence a higher profit for the auction company.

Looking at Figure 6, it can be seen that the throughput offered by both architectural designs is very similar, although the multithreaded architecture shows a slightly better performance when the server is saturated.

**Figure 6** Reply throughput comparison under a dynamic content workload

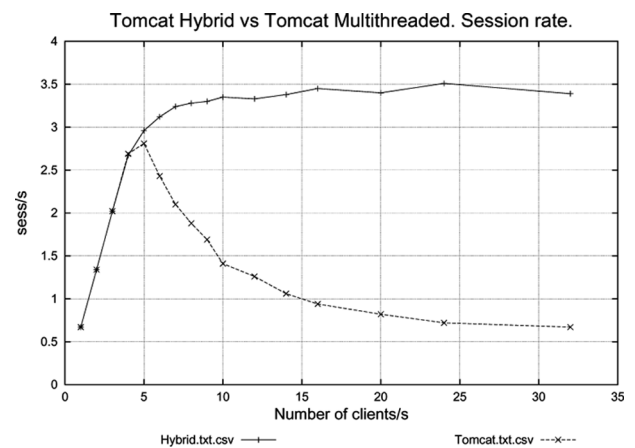


Looking at the number of sessions completed per second, in Figure 7, it can be seen that the amount of successfully finished sessions reached by the hybrid architecture is clearly higher than the amount reached by the multithreaded design, especially beyond saturation. As it can be observed, the multithreaded architecture tends to decrease the number of completed sessions per second as the workload intensity is increased. This can be explained because under a pure multithreaded design the worker threads are obligated to close the client connections in order to be freed and rest available for new incoming connection requests. This situation, under high loads, leads to a scenario where the clients whose connections have been closed by the server, start experiencing problems to be reconnected because the amount of active clients trying to establish a connection is remarkably higher than the amount of worker threads disposed in the server. If this is extended to the amount of connections required to complete a long user session, characteristic of dynamic applications, the probability of being able to finish the session successfully is reduced to near zero, as it has already been discussed in Section 3.

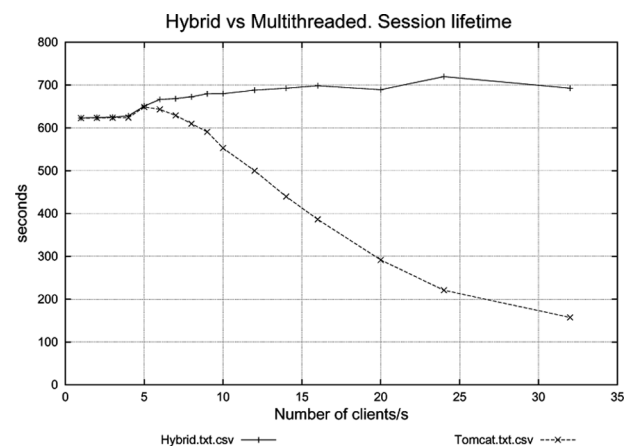
This explanation to the difference of performance between the two architectural designs in terms of session completion is supported by the results shown in Figure 8 that indicate that the sessions completed by

the multithreaded server are significantly shorter than the sessions completed by the hybrid server. This explains that the reply rate for the multithreaded server can be sustained even when the session rate is remarkably reduced because it proves that the sessions completed by the multithreaded server are those ones with less requests (the shorter ones). Completing only short sessions means that many active clients finishes their sessions unsuccessfully, which in turn may result in an important amount of unsatisfied clients that have received a very poor quality of service.

**Figure 7** Successfully completed session rate under a dynamic content workload




**Figure 8** Life time comparison for the sessions completed successfully under a dynamic content workload



Finally, another benefit of the hybrid architecture can be seen in Figure 9, where the amount of requests that have not produced a response before a client timeout expires is represented. The client timeout was set to 10s for the experiment (which is in the limit that a human client usually waits before retrying a request or even discarding it), so the figure represents the amount of times a client has not received a reply for a request 10s after having sent it. As it can be seen, the value for the hybrid server is clearly lower (and better) than the value offered by the multithreaded server. It can be explained because the number of concurrent threads inside the servlets container (see Section 3 for a detailed discussion) is lower under a hybrid server design, which in turn results in

a lower contention and a better performance for each individual request. A multithreaded web server must be configured with a higher number of worker threads in order to support high loads with an acceptable session completion rate but it also leads to a concurrency level that can not be supported in the servlets container, with the result of unacceptable response times for the individual replies.

**Figure 9** Number of client timeouts under a dynamic content workload



## 7 Conclusions and future work

In this paper we have shown how a web server hybrid architecture that combines the best of a multithreaded design with the best of an event-driven model can be used to obtain a high performance web server, especially when the throughput is measured in successfully completed user sessions per second. The proposed implementation into the Tomcat 5.5 code offers a slightly better performance than the original multithreaded Tomcat server when it is tested for a static content application, and a remarkable performance increase when it is compared for a dynamic content scenario, where each user session failure can be put into relation with business revenue losses. Additionally, the natural way of programming introduced by the multithreading paradigm can be maintained for most of the web container code.

Further research will be done toward the exploitation of the hybrid architecture benefits in the area of session based admission control, especially in presence of secure connections (SSL), where the cost of client reconnections can result in an enormous performance impact (see Guitart et al., 2005, for more details on this topic). The hybrid architecture also reduces the complexity of the tuning of a web container, which is an important step toward the implementation of autonomic servers (IBM Research).

## Acknowledgements

This work is supported by the Ministry of Science and Technology of Spain and the European Union

(FEDER funds) under contract TIN2007-60625. Please visit the <http://www.bsc.es/autonomic> for additional information.

## References

- Amza, C., Chanda, A., Cecchet, E., Cox, A., Elnikety, S., Gil, R., Marguerite, J., Rajamani, K. and Zwaenepoel, W. (2002) 'Specification and implementation of dynamic web site benchmarks', *5th IEEE Workshop on Workload Characterization*, Austin, Texas, USA, pp.3–13.
- Barford, P. and Crovella, M. (1998) 'Generating representative web workloads for network and server performance evaluation', *Measurement and Modeling of Computer Systems*, pp.151–160.
- Beltran, V., Carrera, D., Torres, J. and Ayguadé, E. (2004) 'Evaluating the scalability of java event-driven web servers', *2004 International Conference on Parallel Processing (ICPP'04)*, pp.134–142.
- Chen, H. and Mohapatra, P. (2002) 'Session-based overload control in qos-aware web servers', *INFOCOM*, New York, NY, USA, Vol. 2, June, pp.516–524.
- Guitart, J., Beltran, V., Carrera, D., Torres, J. and Ayguadé, E. (2005) 'Characterizing secure dynamic web applications scalability', *19th International Parallel and Distributed Symposium (IPDPS'05)*, IEEE Computer Society, Washington, DC, USA, p.108.1, ISBN: 0-7695-2312-9, DOI: <http://dx.doi.org/10.1109/IPDPS.2005.137>.
- Hu, J., Pyarali, I. and Schmidt, D. (1998) 'Applying the proactor pattern to high-performance web servers', *Proceedings of the 10th International Conference on Parallel and Distributed Computing and Systems IASTED*, October, Las Vegas, Nevada, USA.
- Hu, J. and Schmidt, D. (1999) *Domain-Specific Application Frameworks: Frameworks Experience by Industry*, Wiley & Sons.
- Keller, A. and Ludwig, H. (2002) 'Defining and monitoring service-level agreements for dynamic e-business', *LISA*, pp.189–204.
- Mosberger, D. and Jin, T. (1998) 'httperf: A tool for measuring web server performance', *First Workshop on Internet Server Performance*, ACM, June, pp.59–67.
- Pai, V.S., Druschel, P. and Zwaenepoel, W. (1999) 'Flash: an efficient and portable web server', *Proceedings of the USENIX 1999 Annual Technical Conference*, USENIX Association, Monterey, California, Berkeley, CA, USA, pp.15–15.
- Welsh, M., Culler, D.E. and Brewer, E.A. (2001) 'SEDA: an architecture for well-conditioned, scalable internet services', *Symposium on Operating Systems Principles*, pp.230–243.

## Websites

- IBM Research, *Autonomic Computing*, See <http://www.research.ibm.com/autonomic>
- Jakarta Project. Apache Software Foundation, *Tomcat*, See <http://jakarta.apache.org/tomcat>