



Scaling Webapps With Servlet 3.1 Async I/O



- **Simone Bordet**

- sbordet@intalio.com
- @simonebordet

- **Open Source Contributor**

- Jetty, CometD, MX4J, Foxtrot, LiveTribe, JBoss, Larex

- **Lead Architect at Intalio/Webtide**

- Jetty's HTTP/2, SPDY and HTTP client maintainer


- **CometD project leader**


- Web messaging framework

- **JVM tuning expert**


- **Why Async ?**
- **Async Everywhere**
- **Async I/O API in the Servlet 3.1 Specification**
- **Correctly use the new Servlet Async I/O APIs**

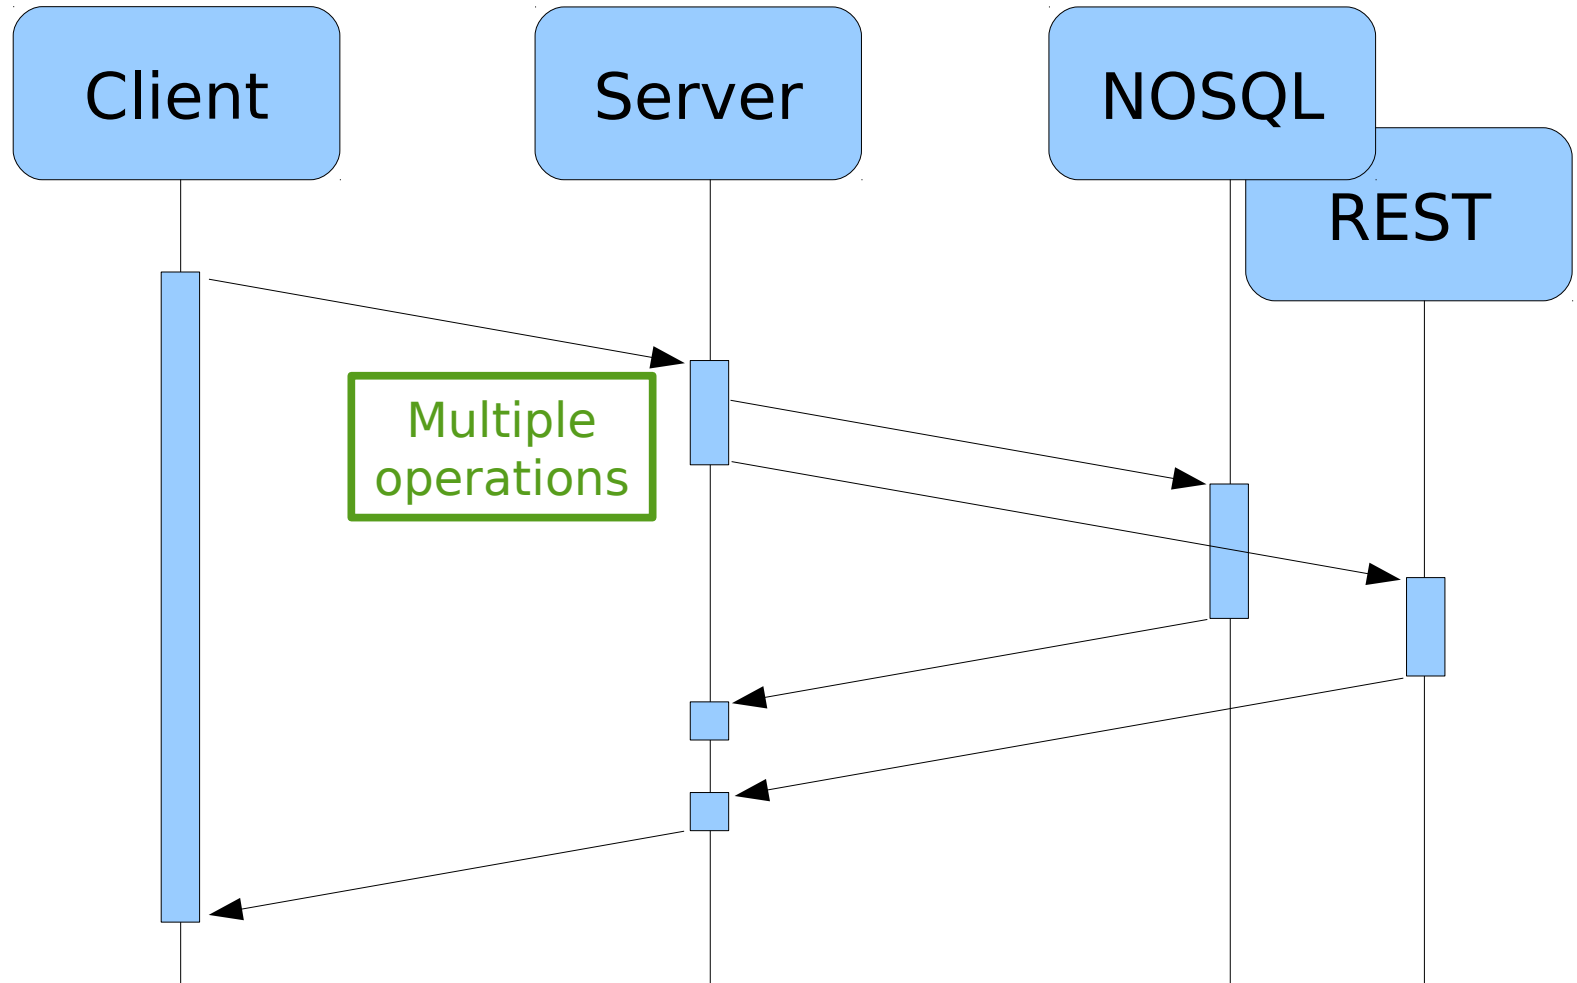
Why Async ?








- **Why blocking waits are bad ?**
- **They consume resources**
 - Native Thread
 - Native Memory (thread stack: 1 MiB per thread)
 - OS scheduler data structures
 - ThreadLocals
 - Local variables in the stack frames
 - GC has to walk the stack frames
- **Blocking waits == \$\$\$**







- **Improved latency**
 - By performing tasks concurrently with less resources
- **Better resource utilization**
 - The SAME thread can serve multiple clients
- **Async waits == \$**
- **Async == Increased performance**
- **Increased performance == \$\$\$**



https://webtide.com/async-rest/

The screenshot shows a web browser window with the URL `https://webtide.com/async-rest/`. The page content is divided into four quadrants, each showing performance data for a specific operation. The top-left quadrant shows blocking performance for 'kayak', with a total time of 199.6ms and a thread held for 199.6ms, accompanied by a red bar. The top-right quadrant shows blocking performance for 'mouse,beer,gnome', with a total time of 593.8ms and a thread held for 593.8ms, accompanied by a red bar. The bottom-left quadrant shows asynchronous performance for 'kayak', with a total time of 200.4ms, a thread held for 1.0ms (0.9 initial + 0.1 generate), and an async wait of 199.5ms, accompanied by a green bar. The bottom-right quadrant shows asynchronous performance for 'mouse,beer,gnome', with a total time of 279.6ms, a thread held for 1.2ms (1.1 initial + 0.2 generate), and an async wait of 278.4ms, accompanied by a green bar. Each quadrant also features a horizontal row of small icons representing different assets.

Operation	Blocking Total Time	Blocking Thread Held	Asynchronous Total Time	Asynchronous Thread Held	Asynchronous Async Wait
kayak	199.6ms	199.6ms	200.4ms	1.0ms (0.9 initial + 0.1 generate)	199.5ms
mouse,beer,gnome	593.8ms	593.8ms	279.6ms	1.2ms (1.1 initial + 0.2 generate)	278.4ms

Async Everywhere

- **Until few years ago parallelism was not widespread**
 - Single core CPUs
- **Modern hardware has gone multi-core**
 - We want to be able to use those cores
- **Future hardware will have even more cores**
- **We need to leverage hw parallelism in software**

- **Async programming is not new**
 - Spawning a thread is a primitive in OS since day one
- **Async programming is difficult...**
 - ... to write: callback hell, synchronization
 - ... to read: our brain reads from top to bottom
- **New APIs are born**
 - epoll, promises, actors, channels, etc.
- **New products are born**
 - Nginx vs Apache, etc.

■ Example: Java 8 CompletableFuture

CompletableFuture

```
.supplyAsync(() -> fetchQueryFromREST(uri))  
.thenApplyAsync(sql -> fetchDataFromNOSQL(sql))  
.thenAccept(data -> updateUI(data));
```

■ We want a way to read code as if it is sequential

- But in fact it is asynchronous

■ You have to marry the API model

- Always use the model, no exceptions

■ Servlet 3.0: async processing of response

```
void doGet(request, response)
{
    OutputStream out = response.getOutputStream();
    AsyncContext ctx = request.startAsync();
    doAsyncREST(request).thenAccept(json -> {
        out.write(json);
        ctx.complete();
    });
}
```

- **Two requirements for an async Servlet**
- **First requirement: `doAsyncREST()`**
 - Must use async libraries to do its work
 - Must return a `CompletableFuture`
 - Or take a callback parameter (or something similar)
- **Second requirement: `thenAccept()` lambda**
 - Must use async libraries

- **X JDBC**
- **X InputStream - OutputStream**
- **✓ Asynchronous I/O (NIO2: files, sockets)**
- **✓ REST**
- **✓ Call other Servlets**

■ Servlet 3.0: no async I/O

```
void doGet(request, response)
{
    OutputStream out = response.getOutputStream();
    AsyncContext ctx = request.startAsync();
    doAsyncREST(request).thenAccept(json -> {
        out.write(json); Blocking!
        ctx.complete();
    });
}
```

- Finally! That's why Servlet 3.1 introduced async I/O
- What API to use then ?
- Use `java.util.concurrent.Future` ?

```
Future<Integer> f = out.write(json);  
int written = f.get().intValue();
```

- Finally! That's why Servlet 3.1 introduced async I/O !
- What API to use then ?
- Use `java.util.concurrent.Future` ?

```
Future<Integer> f = channel.write(json);  
int written = f.get().intValue();
```

**This is just delayed
blocking, not async**

■ Use NIO2 ?

```
out.write(json, null, new CompletionHandler()
{
    void completed(Integer written, Object obj)
    {
        ctx.complete();
    }

    void failed(Throwable x, Object obj) {
    }
}
```


■ Use NIO2 ?

Allocation

Not a lambda

```
out.write(json, null, new CompletionHandler()
```

```
{ Not thread efficient
```

```
void completed(Integer written, Object obj)
```

```
{  
    ctx.complete();  
}
```

```
Prone to stack overflow  
(resolved by further  
thread dispatching)
```

```
void failed(Throwable x, Object obj) {
```

```
}
```

```
}
```

Servlet 3.1 Async I/O

■ Servlet 3.1 async write I/O APIs

```
void doGet(request, response) {
    ServletOutputStream out = response.getOutputStream();
    AsyncContext ctx = request.startAsync();
    out.setWriteListener(new WriteListener() {
        void onWritePossible() {
            while (out.isReady()) {
                byte[] buffer = readFromSomewhere();
                if (buffer != null)
                    out.write(buffer);
            }
            else
                ctx.complete();
        }
    });
}
```

■ Servlet 3.1 async write I/O APIs

```
void doGet(request, response) {
    ServletOutputStream out = response.getOutputStream();
    AsyncContext ctx = request.startAsync();
    out.setWriteListener(new WriteListener() {
        void onWritePossible() {
            Iterative while (out.isReady()) {
                byte[] buffer = readFromSomewhere();
                if (buffer != null)
                    out.write(buffer); Async write
                else
                    ctx.complete(); break;
            }
        }
    });
}
```

■ Normal writes (fast connection)

- Iterative writes in the same thread
- No thread dispatch
- No stack overflow
- No allocation

■ Incomplete writes (slow connection)

- Some iterative write, then `isReady()` returns false
- When ready again, new thread calls `onWritePossible()`

■ Got it ! Is it really this simple ?

■ Bug #1

```
void onWritePossible()  
{  
    out.write(allContent);  
    logger.debug("written: {}", out.isReady());  
}
```

■ Bug #1

```
void onWritePossible()  
{  
    out.write(allContent);  
    logger.debug("written: {}", out.isReady());  
}
```

When `isReady() == false`
`onWritePossible()` will
be called again !

■ Content is written twice !

■ Bug #2

```
void onWritePossible()
{
    if (dataAvailable() && out.isReady())
        out.write(getData());
}
```


■ Bug #2

```
void onWritePossible()  
{  
    if (dataAvailable() && out.isReady())  
        out.write(getData());  
}
```

The order of the expressions in the if statement is significant !

- If `dataAvailable()` returns false, `isReady()` is not called and `onWritePossible()` will never be called
- Switching the if expressions changes the semantic

■ Bug #3

```
void onWritePossible()
{
    if (out.isReady()) {
        out.write("<h1>Async</h1>");
        out.write("<p>Content</p>");
    }
}
```

■ Bug #3

```
void onWritePossible()
{
    if (out.isReady()) {
        out.write("<h1>Async</h1>");
        out.write("<p>Content</p>");
    }
}
```

The first write
may be pending !

- You **HAVE** to use `isReady()` after each write

■ Bug #4

```
void onWritePossible()
{
    if (out.isReady()) {
        out.write("<h1>Async</h1>");
        if (out.isReady())
            out.write("<p>Content</p>");
    }
}
```

■ Bug #4

```
void onWritePossible()
{
    if (out.isReady()) {
        out.write("<h1>Async</h1>");
        if (out.isReady())
            out.write("<p>Content</p>");
    }
}
```

Content written
multiple times !

■ Don't misuse `isReady()`

■ Bug #5

```
void onWritePossible()
{
    int read = file.read(buffer);
    while (out.isReady()) {
        out.write(buffer);
        if (file.read(buffer) < 0)
            ctx.complete(); return;
    }
}
```

■ Bug #5

```
void onWritePossible()
{
    int read = file.read(buffer);
    while (out.isReady()) {
        out.write(buffer);
        if (file.read(buffer) < 0)
            ctx.complete(); return;
    }
}
```

Write may be pending: you don't own the buffer

■ The buffer is overwritten by the new read

■ Bug #6

```
void onWritePossible()  
{  
    try {  
        ...  
    } catch (Exception x) {  
        response.sendError(500);  
    }  
}
```


■ Bug #6

```
void onWritePossible()  
{  
    try {  
        ...  
    } catch (Exception x) {  
        response.sendError(500);  
    }  
}
```

**sendError() is a
blocking API !**

■ Cannot mix async I/O and blocking I/O

- **When ServletOutputStream is put into async mode:**
 - Quacks like a stream
 - Walks like a stream
 - But it's NOT a stream !
- **Broken semantic:**
 - `OutputStream.write()` is expected to be blocking !
 - `ServletOutputStream.isReady()`: a query with side effects !
- **Breaking semantic consequences:**
 - Filters and wrappers are broken !
 - Cannot tell whether the stream is async or not

■ Servlet 3.1 async read I/O APIs

```
void doPost(request, response)
{
    AsyncContext ctx = request.startAsync();
    ServletInputStream in = request.getInputStream();
    in.setReadListener(new ReadListener() {
        void onDataAvailable() {
            while (in.isReady() && !in.isFinished()) {
                int len = in.read(buffer);
                if (len > 0)
                    store(buffer); // Blocking API
                if (in.isFinished())
                    ctx.complete(); return;
            }
        }
    });
}
```

■ Servlet 3.1 async read I/O APIs

```
void doPost(request, response)
{
    AsyncContext ctx = request.startAsync();
    ServletInputStream in = request.getInputStream();
    in.setReadListener(new ReadListener() {
        void onDataAvailable() {
            Iterative while (in.isReady() && !in.isFinished()) {
                int len = input.read(buffer); Async read
                if (len > 0)
                    store(buffer); // Blocking API
                if (in.isFinished())
                    ctx.complete(); return;
            }
        }
    });
}
```

■ Normal reads (fast connection)

- Iterative reads in the same thread
- No thread dispatch
- No stack overflow
- No allocation

■ Incomplete reads (slow connection)

- Some iterative reads, then `isReady()` returns false
- When ready again, new thread calls `onDataAvailable()`

■ Got it ! Is it really this simple ?

■ Bug #1

```
void onDataAvailable() {  
    while (in.isReady()) {  
        in.read(buffer);  
        store(buffer);  
    }  
}
```

■ Bug #1

```
void onDataAvailable() {  
    while (in.isReady()) {  
        in.read(buffer);  
        store(buffer);  
    }  
}
```

isReady() returns true at EOF

■ Infinite loop

- The loop must call also `in.isFinished()`

■ Bug #2

```
void onDataAvailable() {
    while (in.isReady() && !in.isFinished()) {
        int read = in.read(buffer);
        store(buffer, new Callback() {
            void done() {
                onDataAvailable();
            }
        });
    }
}
```


■ Bug #2

```
void onDataAvailable() {
    while (in.isReady() && !in.isFinished()) {
        int read = in.read(buffer);
        store(buffer, new Callback() {
            void done() {
                onDataAvailable(); Stack overflow !
            }
        });
    }
}
```

■ Stack overflows for fast connections

■ Mixing different async models requires more careful coding

- Jetty's `IteratingCallback` is one solution

■ Bug #3

```
void onDataAvailable() {  
    String id = request.getParameter("id");  
    ...  
}
```

■ Bug #3

```
void onDataAvailable() { Blocking API !  
    String id = request.getParameter("id");  
    ...  
}
```

■ Same with request.getParts()

- Multipart APIs

Echo Servlet Blocking I/O

```
public class SyncServlet extends HttpServlet
{
    protected void service(HttpServletRequest request, HttpServletResponse response) throws ServletException, IOException
    {
        byte[] buffer = new byte[BUFFER_SIZE];
        while (true)
        {
            int read = request.getInputStream().read(buffer, 0, BUFFER_SIZE);
            if (read < 0)
                break;
            response.getOutputStream().write(buffer, 0, read);
        }
    }
}
```

Echo Servlet Async I/O

```
public static class AsyncServlet extends HttpServlet
{
    @Override
    protected void service(HttpServletRequest request, HttpServletResponse response) throws ServletException, IOException
    {
        AsyncContext asyncContext = request.startAsync(request, response);
        asyncContext.setTimeout(0);
        Echoer echoer = new Echoer(asyncContext);
        request.getInputStream().setReadListener(echoer);
        response.getOutputStream().setWriteListener(echoer);
    }

    private class Echoer implements ReadListener, WriteListener
    {
        private final byte[] buffer = new byte[BUFFER_SIZE];
        private final AsyncContext asyncContext;
        private final ServletInputStream input;
        private final ServletOutputStream output;
        private boolean complete;

        private Echoer(AsyncContext asyncContext) throws IOException
        {
            this.asyncContext = asyncContext;
            this.input = asyncContext.getRequest().getInputStream();
            this.output = asyncContext.getResponse().getOutputStream();
        }

        @Override
        public void onDataAvailable() throws IOException
        {
            while (input.isReady())
            {
                int read = input.read(buffer);
                output.write(buffer, 0, read);

                if (!output.isReady())
                    return;
            }

            if (input.isFinished())
            {
                complete = true;
                asyncContext.complete();
            }
        }

        @Override
        public void onAllDataRead() throws IOException
        {
        }

        @Override
        public void onWritePossible() throws IOException
        {
            if (input.isFinished())
            {
                if (!complete)
                    asyncContext.complete();
            }
            else
            {
                onDataAvailable();
            }
        }

        @Override
        public void onError(Throwable failure)
        {
            failure.printStackTrace();
        }
    }
}
```

■ Async I/O Echo Servlet

■ Pros

- Async

■ Cons

- 4x-6x more code
- Requires a state machine (a simple boolean, but still)
- More difficult buffer handling
- Complexity, maintainability, etc.

Conclusions

- **Blocking code has a reason to exist**
 - And will continue to exist

- **Is the async I/O complexity worth it ?**
 - For a class of applications, YES !
 - Improved scalability
 - Decreased latency
 - Makes you more money

<https://webtide.com/async-rest/>




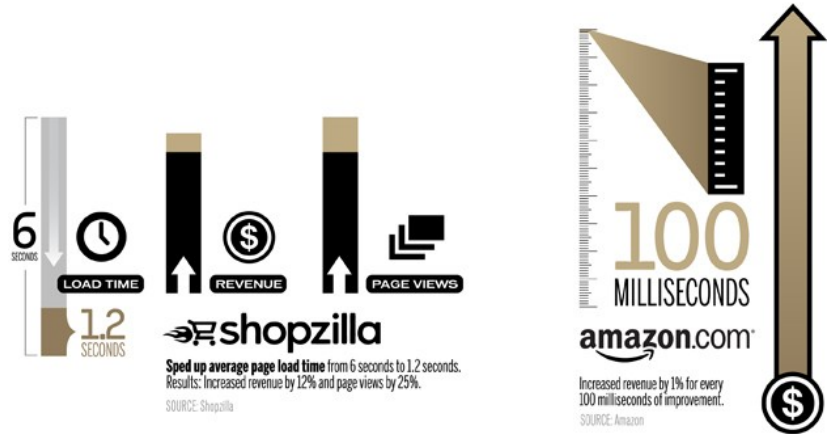
The screenshot shows a web browser window with the URL <https://webtide.com/async-rest/>. The page content is divided into four quadrants, each showing performance metrics for a specific operation. The top-left quadrant shows blocking metrics for 'kayak', with a total time of 199.6ms and a thread held for 199.6ms, accompanied by a red bar. The top-right quadrant shows blocking metrics for 'mouse,beer,gnome', with a total time of 593.8ms and a thread held for 593.8ms, accompanied by a red bar. The bottom-left quadrant shows asynchronous metrics for 'kayak', with a total time of 200.4ms, a thread held for 1.0ms (0.9 initial + 0.1 generate), and an async wait of 199.5ms, accompanied by a green bar. The bottom-right quadrant shows asynchronous metrics for 'mouse,beer,gnome', with a total time of 279.6ms, a thread held for 1.2ms (1.1 initial + 0.2 generate), and an async wait of 278.4ms, accompanied by a green bar. Each quadrant also features a row of small images representing different assets.

Blocking: kayak
Total Time: 199.6ms
Thread held (red): 199.6ms

Blocking: mouse,beer,gnome
Total Time: 593.8ms
Thread held (red): 593.8ms

Asynchronous: kayak
Total Time: 200.4ms
Thread held (red): 1.0ms (0.9 initial + 0.1 generate)
Async wait (green): 199.5ms

Asynchronous: mouse,beer,gnome
Total Time: 279.6ms
Thread held (red): 1.2ms (1.1 initial + 0.2 generate)
Async wait (green): 278.4ms



strangeloop www.strangeloopnetworks.com

Questions & Answers