

## Steamship Authority - January 21 2020

### Summary

During the Martha's Vineyard opening, the infrastructure behind the Steamship Authority website wasn't able to handle the volume of traffic (ie. the number of concurrent visitors to the website). Because traffic passed the limits of multiple parts of the infrastructure, Steamship and Imarc spent the day identifying one problem before they could see what the next bottleneck in the infrastructure was; with Imarc making code changes to the mobile site to connect more efficiently to the Reservation System; changing configuration settings on the Web Database servers; changing the website connections to the database servers to be more efficient; and reducing the allowed connections to each website server.

Imarc has been a vendor for Steamship Authority since 2010 and built the current mobile website and main website. Steamship regularly contracts Imarc to complete specific changes to the Steamship websites. On January 6, Steamship specifically asked for Imarc to be available for each of the openings. Actions Imarc took were all with the approval or direction of the Steamship Authority's team through constant communication throughout the day. The servers involved are hosted and maintained by Steamship.

The Steamship Authority doesn't have any kind of service or support agreement with Imarc. Imarc does not provide support outside of normal business hours.

### Quick Reference of Involved Infrastructure

- The main website, [www.steamshipauthority.com](http://www.steamshipauthority.com), is served from five web servers.
- The web servers communicate with the Reservation System for reservation details.
- The web servers also communicate with a set of two Web Server Database servers for shared files, CMS content, and transient, operational information (like tracking who is logged in.)
- The mobile website, [m.steamshipauthority.com](http://m.steamshipauthority.com), is served from two separate web servers.
- The mobile servers also communicate with the Reservation System for the list of trips and their status, but reservations cannot be booked directly from the mobile website.
- The mobile servers don't need to use a Web Server Database at all.
- The Web Database Service is an application that runs on the web database servers.

### Timeline

4:50 am	Imarc and Steamship meet via conference call. This conference call was used throughout the day to maintain communication between Imarc and Steamship.
5:00	Steamship updates the dates available within the Reservation System and visitors hit the site.
5:01	Reservation System starts sending errors to the web servers. These errors include session disconnects and session not available.

5:05	Web servers start to return some error messages about not being able to connect to the Web Server Database servers.
~5:45	Steamship restarts the Web Server Database servers, but the Web Server Database service failed to restart.
~6:00	Steamship changes the available dates back with the plan to reschedule the opening for Martha's Vineyard to later in the day.
~6:00	<p>Imarc determined that a system limit prevented the Web Server Database from service from starting, but neither Steamship nor Imarc knew how to change this system limit at this time. With Steamship's approval, Imarc configured the Web Server Database to with reduced settings so that it was under the system limit and started the service. At this point, both teams believed that, while reduced, the Web Server Database would be able to keep up with visitor traffic and provide a better experience to customers while both teams continued to troubleshoot.</p> <p>Technical: Imarc limited the number of connections to the Web Server Database to 1,000 because the services would not allow it to start with 2,000+ connections due to a system limit.</p>
~8:30	Opening time was rescheduled to 12:00 p.m.
11:10	Under Steamship's direction, Imarc made a code change to the mobile website to change the mobile site to use a persistent connection – a similar change to what was made on the desktop site in January 2018 – based on the theory that visitor traffic to the mobile site was putting too much strain on the Reservation System. This change reduced the number of connections to the Reservation System immediately, so we felt confident going into the 12:00 p.m. opening.
12:00 p.m.	Second opening time, Steamship updates the dates available.
~12:00	<p>Almost immediately, we see the Web Server Database servers failing to keep up with the web servers and the website go down.</p> <p>Technical: Imarc immediately saw that the web servers were running out of connections to the Web Server Database. At this point, the limit was 1,000 connections.</p>
~12:30	Under Steamship's direction, Imarc changed the main site to use a persistent connection to the Web Server Database service, but it didn't improve performance.
~1:10	<p>Through further research and investigation into Steamship's Web Database servers, Imarc found a way to raise the system limit preventing the Web Server Database from running with more connections.</p> <p>With Steamship's approval, Imarc changed both this system limit for Steamship as well as the number of connections to the Web Server Database service from 1,000 back to 16,000.</p>
~1:30	The web servers stop responding and the website goes down again. Steamship physically restarts them. Imarc believes the servers stopped responding due to the work on the Web Server Database service.

~1:30-2:00	Performance appears normal.
~2:10	Web servers stop responding and the website goes down again. At Imarc's request, Steamship physically cycles them and Imarc goes in to debug and determined the Web Servers were running out of memory.
~2:30	With Steamship's approval, Imarc reduced the allowed connections to each web server from 2,800 to 500 to prevent them from running out of memory and restarted the servers.
3:30	Everything was running smoothly after reducing the allowed connections to each web server. We agree to inform customers the site is operating as expected.

## Issues

### 1. Reservation System Connectivity

The first issue we saw was the Reservation System sending errors to the web servers.

- **Cause:** we believe this issue was caused by load generated by visitors to the mobile site. [SSA to provide metrics]
- **Why wasn't this detected:** Previously, Steamship did not include the mobile site in their load tests as it wasn't believed that the data it was fetching from the Reservation System could be significant enough to interfere with the Reservation System's performance.
- **Fix:** Change the mobile site to use a persistent connection to the Reservation System and include the mobile site in all future load testing.
- **Follow up:** there are a few other performance changes Imarc could duplicate from the main website to the mobile website to further reduce the load it puts on the Reservation System, such as caching and query optimization.

### 2. Web Server Database Not Restarting

After Steamship restarted the Web Server Database servers, the Web Server Database service should have automatically restarted on both servers.

- **Cause:** System settings on the Web Server Database servers prevented the Web Server Database service from starting after the servers were rebooted. More specifically, the Web Server Database service wouldn't start because the configured number of maximum connections, 16,000, was above the system's configured limits.
- **Why wasn't this detected:** Imarc believes that these limits reverted after a previous reboot of the Web Server Database servers and the load on the web servers hadn't been high enough to hit this limit. Load testing has been primarily targeted at the Reservation System performance, and might not stress the Web Server Database service significantly.
- **Fix:** Reconfigured the proper limits of the system configuration files.
- **Follow up:** we should verify that the proper system settings are configured for the given hardware across production servers. Ideally, Imarc would work with Steamship to decide what Steamship would like the limits at (as they're related to load, redundancy, and stability) or whether we should be investigating different solutions (like connection pooling or reducing the amount of data pulled from the Web Server Database.)

While monitoring on January 21, 2020 Imarc saw the number of active connections to the Web Server Database peak around 5,500. For comparison, the current number of connections at 1:00 p.m. on January 22, 2020, was approximately 55.

### Web Server Database Connectivity (Out of Connections)

This is actually just part of the previous issue, but I've separated it out for clarity. Because of the previous issue (Web Database Servers Not Restarting) Imarc configured the Web Server Database with reduced

settings that ultimately couldn't keep up with the website under load. More specifically, Imarc lowered the max connections from 16000 down to 1000 (even at 2000, the service wouldn't start).

- **Cause:** Imarc configured it this way trying to fix the previous issue, but this solution wasn't effective.
- **Fix:** remove this change and go back to the previous configuration once the previous issue was fixed.
- **Follow up:** nothing required.

### 3. Web Servers Not Responding

The website servers stopped responding twice in the early afternoon.

- **Cause:** the website servers ran out of memory under load due to the volume of traffic.
- **Why wasn't this detected:** The previous stress tests that have been run focused on the Reservation System primarily, resulting in the website servers not being stressed heavily.
- **Fix:** the website servers were reconfigured with lower limits (500 per server) on the number of simultaneous connections they'll attempt to handle.
- **Follow up:** reducing the number of simultaneous connections may affect the benchmarks when going through future load testing. The load test should be re-run now and the limits should be reviewed until all parties are happy with the results.

## Mid to Long Term Recommendations

Use Google Analytics across the site – there are ways to anonymize user data while providing benchmark data in order to make better decisions.

Implement a unified log and resource monitoring application across all servers. Tools are available that can provide instant insight into issues across all servers instead of manually browsing through error emails and logs.

Plan for a back-end and front-end upgrade. The typical lifecycle of a website is 4-7 years. While Imarc does not recommend completely rewriting all the code, the core foundation of the site as well as the infrastructure and architecture design dates back to 2012/2013. While the site was developed to handle traffic during that time, as well as significant margin for growth, Imarc believes the website has been hitting the limits of this 7+ year old foundation. While Imarc has made some significant changes for performance over the years, the lifespan of the website should be considered. The Steamship site has a lot of physical resources available, but there are some fundamental code and architecture updates that Imarc recommends. Examples include more robust caching, connection pooling, utilizing a network share, and rearchitecting the reservation process.