

## BlueBRIDGE - Support #11323

VRE # 10195 (Released): FAO Training VRE on Stock monitoring tools

### FAO StockMonitoringTools for SDG 14.4.1 - Concurrency problems

Mar 01, 2018 11:56 AM - Enrico Anello

<b>Status:</b>	Closed	<b>Start date:</b>	Mar 05, 2018
<b>Priority:</b>	High	<b>Due date:</b>	
<b>Assignee:</b>	Andrea Dell'Amico	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Sprint:</b>	Computational Service		
<b>Infrastructure:</b>	Development, Pre-Production, Production		
<b>Description</b>			
<p>Hi all, FAO is making an e-learning course for the SDG 14.4.1 and there will be many links to the StockMonitoringTools shiny application embedded in the SDG 14.4.1 VRE.</p> <p>The shiny applications holds a series of interactive functionalities but it seems that we can run them just one at the time; the other sessions gets in hold until the first on the chain has finished and also the ones down in the queue get a timeout after a certain while.</p> <p>Is it possible to increase the number of concurrent sessions to something like 10? Do we need extra power for that like new machines? If yes how much would be the costs?</p> <p>Another solution would be to translate all these methods to the data miner but we would indeed loose all the interactivity. The methods are very light weight for the datasets we will provide to the e-learning course and for this reason we embedded them in the shiny app directly rather than using the data-miner.</p> <p>Can you please advise if it's possible to increase the concurrency of the Shiny Proxy for the StockMonitoringTools application or if we have to get another approach to let multiple users work on this application at the same time?</p>			
<b>Subtasks:</b>			
Support # 11356: Stock Monitoring Tools not working in VRE			<b>Closed</b>

#### History

##### #1 - Mar 01, 2018 12:12 PM - Andrea Dell'Amico

That behaviour was also noted when the app was deployed, but there's no limit on the concurrent instances in the shinyproxy configuration or in the docker environment. I also do not see a memory shortage.

##### #2 - Mar 01, 2018 12:20 PM - Enrico Anello

@[andrea.dellamico@isti.cnr.it](mailto:andrea.dellamico@isti.cnr.it) thanks,

Unfortunately I cannot test the concurrency on my dev environment because the free shiny server I am using does not allow any concurrent session by factory.

I think that the problem might be under the hoods, in particular on the R session each session relies to. Can you please check, somehow, if each session generates a different R session?

Enrico

##### #3 - Mar 01, 2018 04:25 PM - Pasquale Pagano

- Assignee changed from Pasquale Pagano to Andrea Dell'Amico

@[enrico.anello@fao.org](mailto:enrico.anello@fao.org), we do not recommend shinyproxy even if we can understand why you are using it. We do not recommend it because we do not have expertise on it.

@[andrea.dellamico@isti.cnr.it](mailto:andrea.dellamico@isti.cnr.it), does it make sense to replicate the shinyproxy on 2-3 servers and proxy the requests to it?

**#4 - Mar 01, 2018 04:49 PM - Enrico Anello**

thanks @[pasquale.pagano@isti.cnr.it](mailto:pasquale.pagano@isti.cnr.it)

I think the only way we currently have to deploy shiny apps within BlueBridge context is the shiny proxy, isn't it?

**#5 - Mar 01, 2018 08:17 PM - Andrea Dell'Amico**

- Status changed from New to In Progress

Enrico Anello wrote:

I think the only way we currently have to deploy shiny apps within BlueBridge context is the shiny proxy, isn't it?

Yes, it is.

I think that the problem might be under the hoods, in particular on the R session each session relies to. Can you please check, somehow, if each session generates a different R session?

Opening more than one app instance in different browser tabs? Which algorithms should I try? Elefan with the sample dataset?

**#6 - Mar 01, 2018 08:25 PM - Enrico Anello**

@[andrea.dellamico@isti.cnr.it](mailto:andrea.dellamico@isti.cnr.it)

You can try with the Elefan SA algorithm with the sample dataset you can download within the shiny app.

The sample dataset works with all the Elefan methods. Elefan GA takes a few seconds to complete, Elefan SA around a minute while Elefan around 5 minutes.

Cheers,  
Enrico

**#7 - Mar 01, 2018 08:33 PM - Andrea Dell'Amico**

In the meantime I run the Elefan algorithm. From what I saw in the docker logs, the computation starts immediately on every instance (I started the second three minutes after the first, the calculation on the second finished two minutes after the first). I'm going to run more parallel sessions.

**#8 - Mar 02, 2018 12:27 PM - Andrea Dell'Amico**

- % Done changed from 0 to 50

I've found the problem. The server has 4 CPUs, so at most 4 containers are spawned. All the requests after the fourth are queued. So it's a matter of adding more vCPU to the VM and restart it.

Tell me when it can be done.

**#9 - Mar 02, 2018 12:31 PM - Enrico Anello**

@[andrea.dellamico@isti.cnr.it](mailto:andrea.dellamico@isti.cnr.it) great! You can do it even now!  
How many CPUs can be added to the instance?

**#10 - Mar 02, 2018 12:34 PM - Andrea Dell'Amico**

I can rise them to 12. Is it enough?

**#11 - Mar 02, 2018 12:41 PM - Enrico Anello**

should be good for the time being. In the future there might be the need to increase them to 24.

**#12 - Mar 02, 2018 12:48 PM - Andrea Dell'Amico**

- % Done changed from 50 to 100

- Status changed from In Progress to Feedback

We are not able to provide that many resources on a single host, we will have to split the service on more than one VM. That should be completely transparent to the end users, but we'll need some days to build the configuration.

The server is running with 12 CPUs, in the meantime.

**#13 - Mar 02, 2018 12:54 PM - Enrico Anello**

No problem if you need some days for the configuration. The important thing is that we can add more CPUs if needed :)

Thanks @[andrea.dellamico@isti.cnr.it](mailto:andrea.dellamico@isti.cnr.it)

**#14 - Mar 02, 2018 12:59 PM - Andrea Dell'Amico**

- Status changed from Feedback to Closed