

Theme: Shanghai library FOLIO project

Time: March 9, 2021 07:00pm (EST) / March 10, 2020 08:00am (GMT+8)

Attendees:

Vincent Bureau (Enterprise Architect, EBSCO)
Gang Zhou (Project manager, Shanghai library)
Sha Jiang (Technical Director, Jiata)
Lucy Liu (Product Owner, Folio China)

Notes

1. System operations and maintenance in the production environment

Sha Jiang: Troubleshooting is problematic in the production environment. We found it insufficient to troubleshoot simply by analyzing the logs. Server, memory and CPU uses are also needed. How to handle this situation?

Vince:

- The logs are the big source for troubleshooting.
- Lucy posted the question to Slack. Jason's response:



Jason Root 2 days ago

It's a good idea to have some kind of log shipper and indexer, so you can search container logs more easily from all of the modules. We use Rancher's integration with K8s and fluentd to ship our logs off to a Splunk indexer. (edited)



Jason Root 2 days ago

In our experience, it's also a good idea to examine the Postgres database logs, as well as Okapi's logs when issues arise.

We do things similar at EBSCO. We collect logs from all the systems and put them in a central searchable database. To support our reselling partners, we also gather logs from the systems we host for them and make them available through an Elasticsearch database.

- Logs might not be enough when it comes to the operational things. Server, memory and CPU utilization can be useful for troubleshooting. At EBSCO, we run our hosting on Amazon Web Services (AWS) in the cloud. So we have a certain amount of metrics information that is made available to us through that service. We put all of our different systems in various clusters and use the dashboards that AWS provides to give us information on CPU utilization across everything. We don't make that information available to the customers directly or to partners. But we will use those for troubleshooting. I think the most use to get from that is to be proactive to anticipate when there might be a problem so that you can compensate before the customer can see the problem. If you see that there's a rise in memory usage or CPU utilization, you know you might take that opportunity to move some of your clients and usage onto different clusters or maybe a larger capacity situation. So they don't experience the problem.

- Are you hosting by yourself or somewhere else in China?

Sha Jiang: We are hosting with Shanghai Library. It has an IT department with a data center of its own.

Vince: Are you working closely with them or do you have the problem of accessing it? At EBSCO, we have some of our stuff also hosted on premise. We have a couple of data centers. It can be problematic to get access to the information. People who run the data center are not always readily available to provide you with the logs that are diagnostic, or you have limitations to what you can access.

Sha Jiang: We work closely with the operation team. That's not a big problem. If we can't get certain information directly from the server, we can ask someone to get it for us.

Vince:

- That's going to be an obstacle. You need to have some ability to integrate something into those servers to export information. So it is available to you to troubleshoot. That would involve some sort of log search, maybe elasticsearch, something that happens automatically. So you don't have to ask for the logs. But the logs will be there immediately as soon as you need to look at them.
- Similarly, depending on how the IT department has implemented their hosting, I'm assuming that they might be using some kind of virtualization in their hosting to run virtual servers on actual hardware. Generally, the virtualization service will provide you with metrics like CPU utilization and memory usage through dashboards. So you have access to that. And that gives you the ability to analyze what is the problem.
- Folio does not have a centralized mechanism that's very good for collecting and observing all that kind of thing. There were a couple proposals that were made within the Folio project to provide that sort of a central logging and traceability system. But it was never prioritized enough to become part of the community project. The current design involves Okapi, which is a centralized processing Gateway, so it is sort of a single point that everything has to go through. But if we could at some point evolve the platform differently, we could eventually introduce a service mesh type of operational configuration. We would have different services that deployed and wrapped by a proxy around each of them. And the proxy, the envelope, around each of the modules would provide metrics that could be captured in the central location that's part of the service mesh. It would not happen soon. But that's the direction.
- So I do think that the right direction is to aggregate the logs, provide a mechanism to quickly search through them, and at the same time, find a way to expose the dashboards that will probably be available through the virtualization engine that you have.
- It is noticed in some cases that the logs themselves can be the cause of the problem, not the solution to the problem. If your system is writing too many logs into a log capturing environment that is not fast enough, your system will be slowed down by the log writing mechanism. So you should realize that you need a mechanism in place to purge the logs, to move them to another location so you don't constantly accumulate them because you

will also be running into space problems. If you have a physical storage and the log starts to use up all your physical storage room, it will cause additional performance issues.

- You may take a look at the Performance Task Force's wiki pages <https://wiki.folio.org/pages/viewpage.action?spaceKey=FOLIJET&title=Performance+Task+Force+Team>. There is plenty of information on it. You can use carrier.io to diagnose what caused the lag. (Under "Tutorials")
- In terms of the operations, you should take advantage of any kind of Health checking that the modules can provide. They don't tell you much. But they'll tell you if the module is still responding and still working. You probably want to set up some kind of monitoring system that confirms that everything is up and running and potentially send out alerts if you have something that is not up.

2. A friendly reminder that daylight saving time 2021 begins at 2 a.m. March 14. The time difference between Shanghai and the US will change.