Deploying a multi-region, highly available MySQL architecture

CASE STUDY

November 2nd, 2017

Pythian

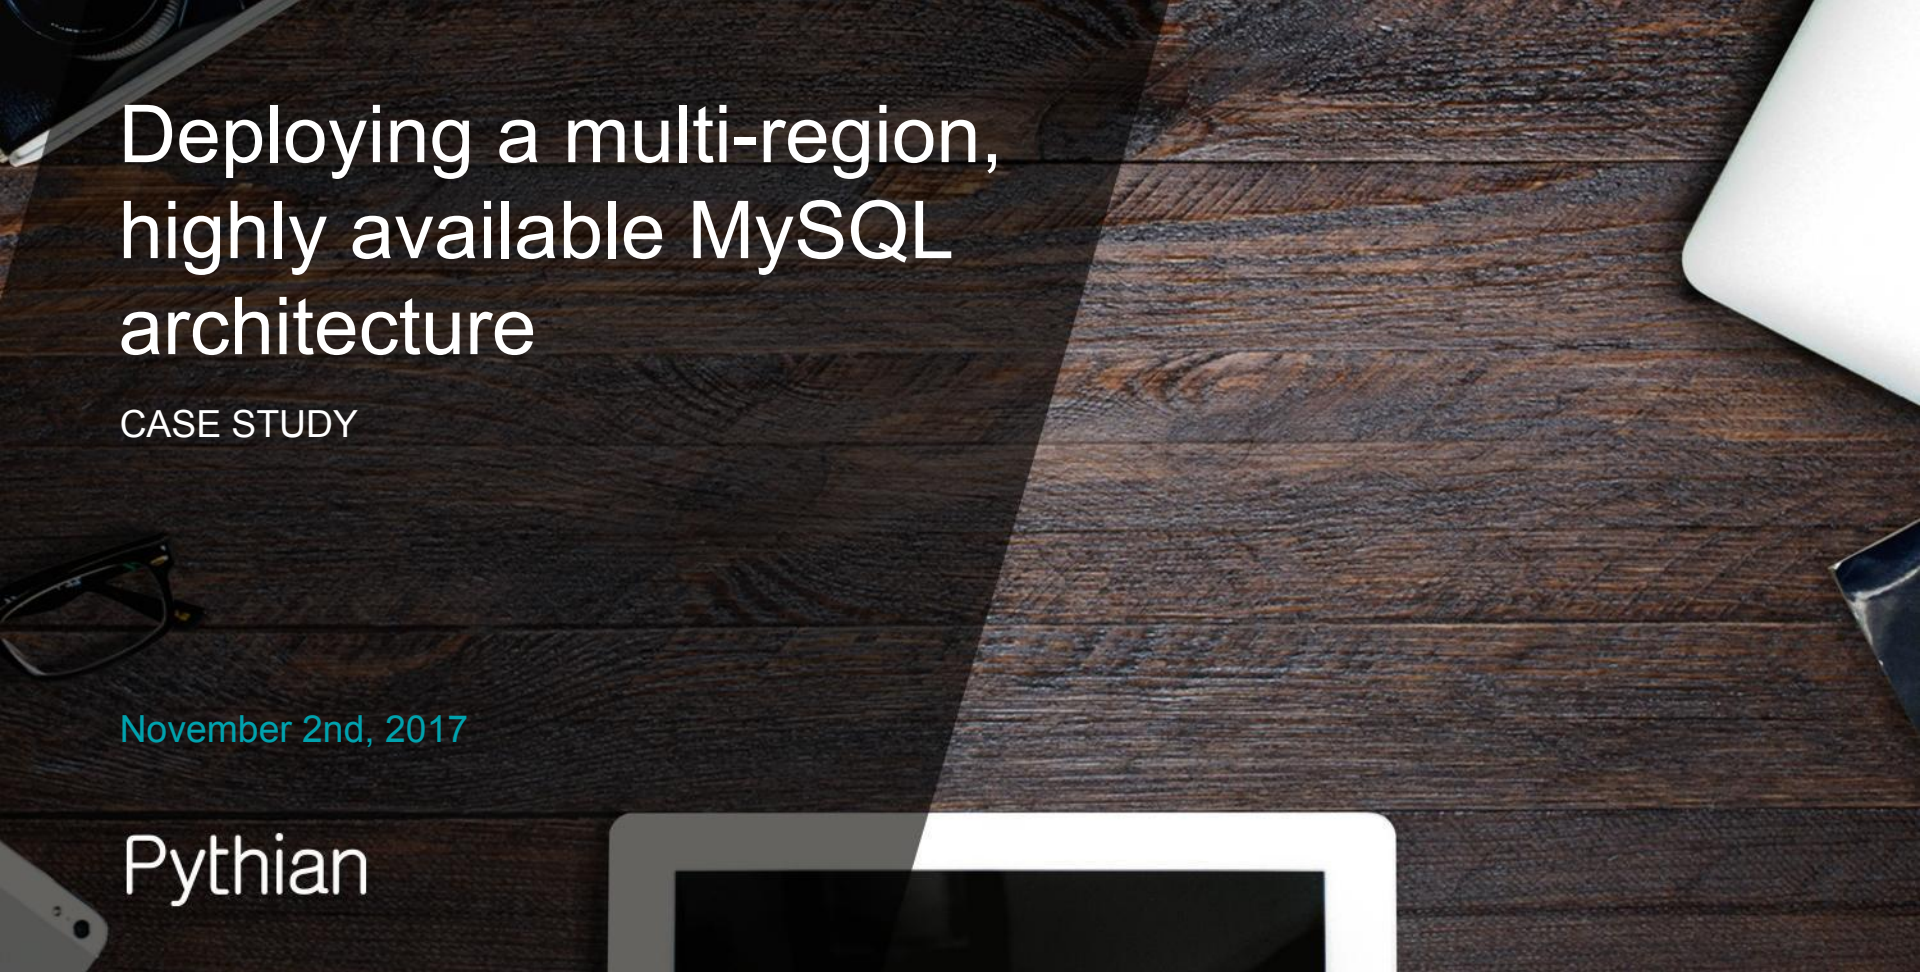

About me..

- OSDB Internal Principal Consultant @Pythian
- Automation super fan
- Working on my master’s thesis (ML - Groups discovery)
- Husband, father (a boy and a girl)
- Based out of Buenos Aires, Argentina, the land of

Nope, that’s Brazil

We are down here

AGENDA

- The requirements
- The Options
- Final architecture
- Failure scenarios
- Limitations
- Q&A

The requirements

The requirements

- Should be in Azure
- Three regions for geo locality
- All data needs to be available from all regions
- MySQL server as the database backend
- Handle failures automatically

Considerations

- Interactive application
- Workload: Mostly reads

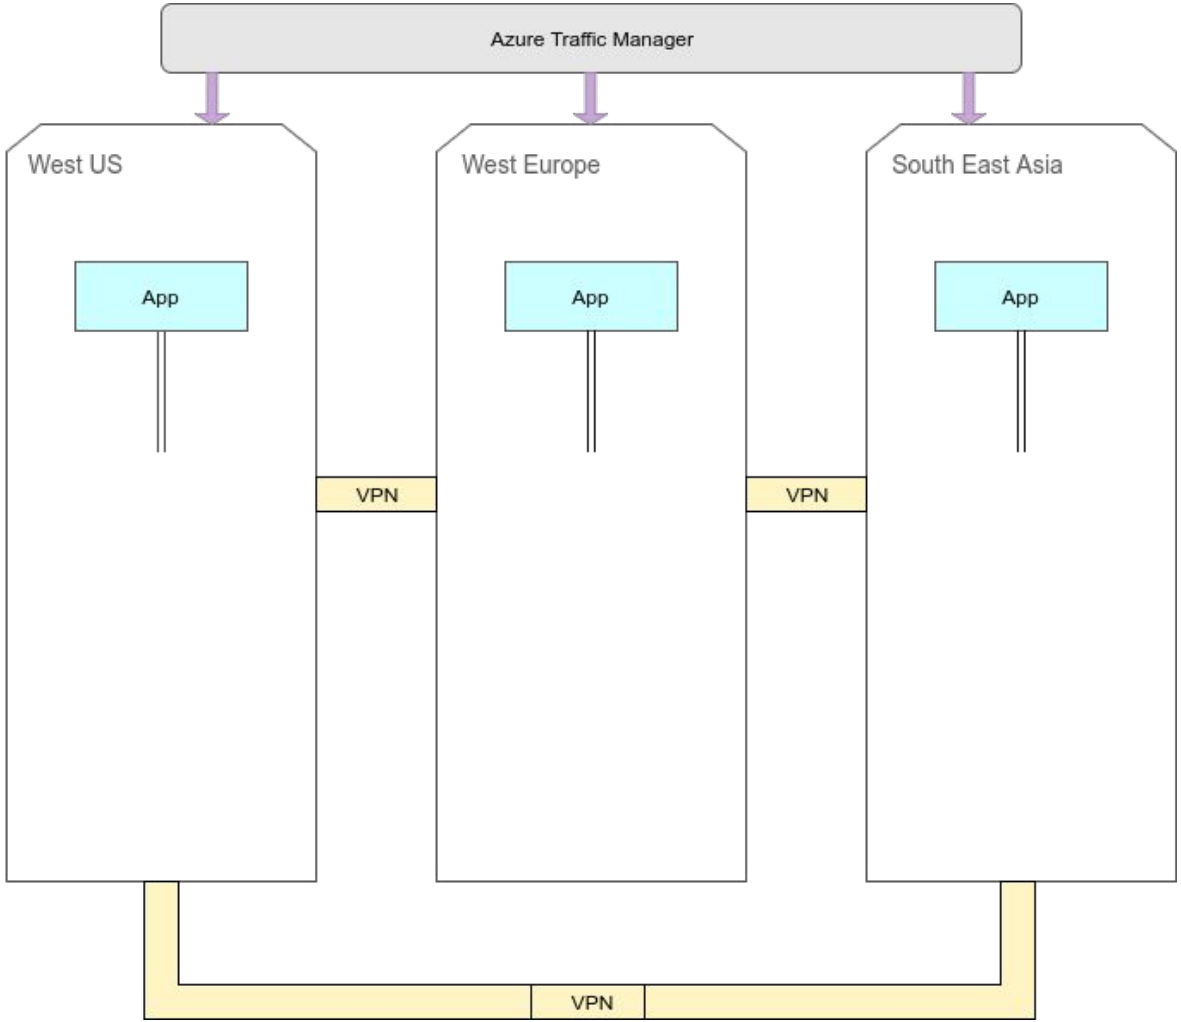

The Options

Data replication

- Multi master
  - Virtually synchronous -> wsrep
- Single master
  - Semi-synchronous
  - Asynchronous

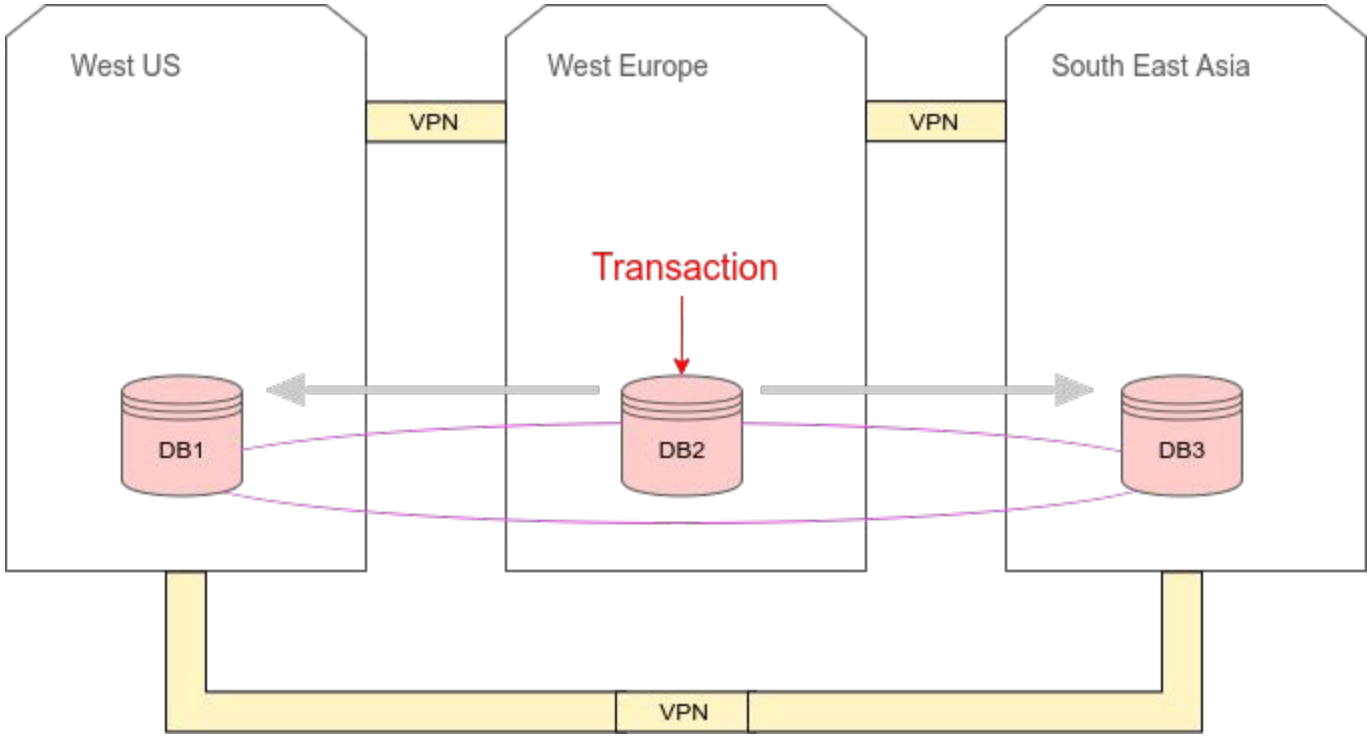

Database traffic routing

- Read-write splitting
- Failures handling
- Proxy placement
  - At the client level
  - As a middle layer
  - At the database server
- Product options
  - MaxScale
  - ProxySQL

ProxySQL configuration

- Read-write splitting with failures handling
- Read after write operations -> /* SQL comment */
- Query caching

Topology manager

- Rebuild replication topology upon master failure
- Fencing
- Topology monitor location
- Product options
  - MHA
  - Github Orchestrator
    - Was not evaluated but it could had been an option

Final architecture

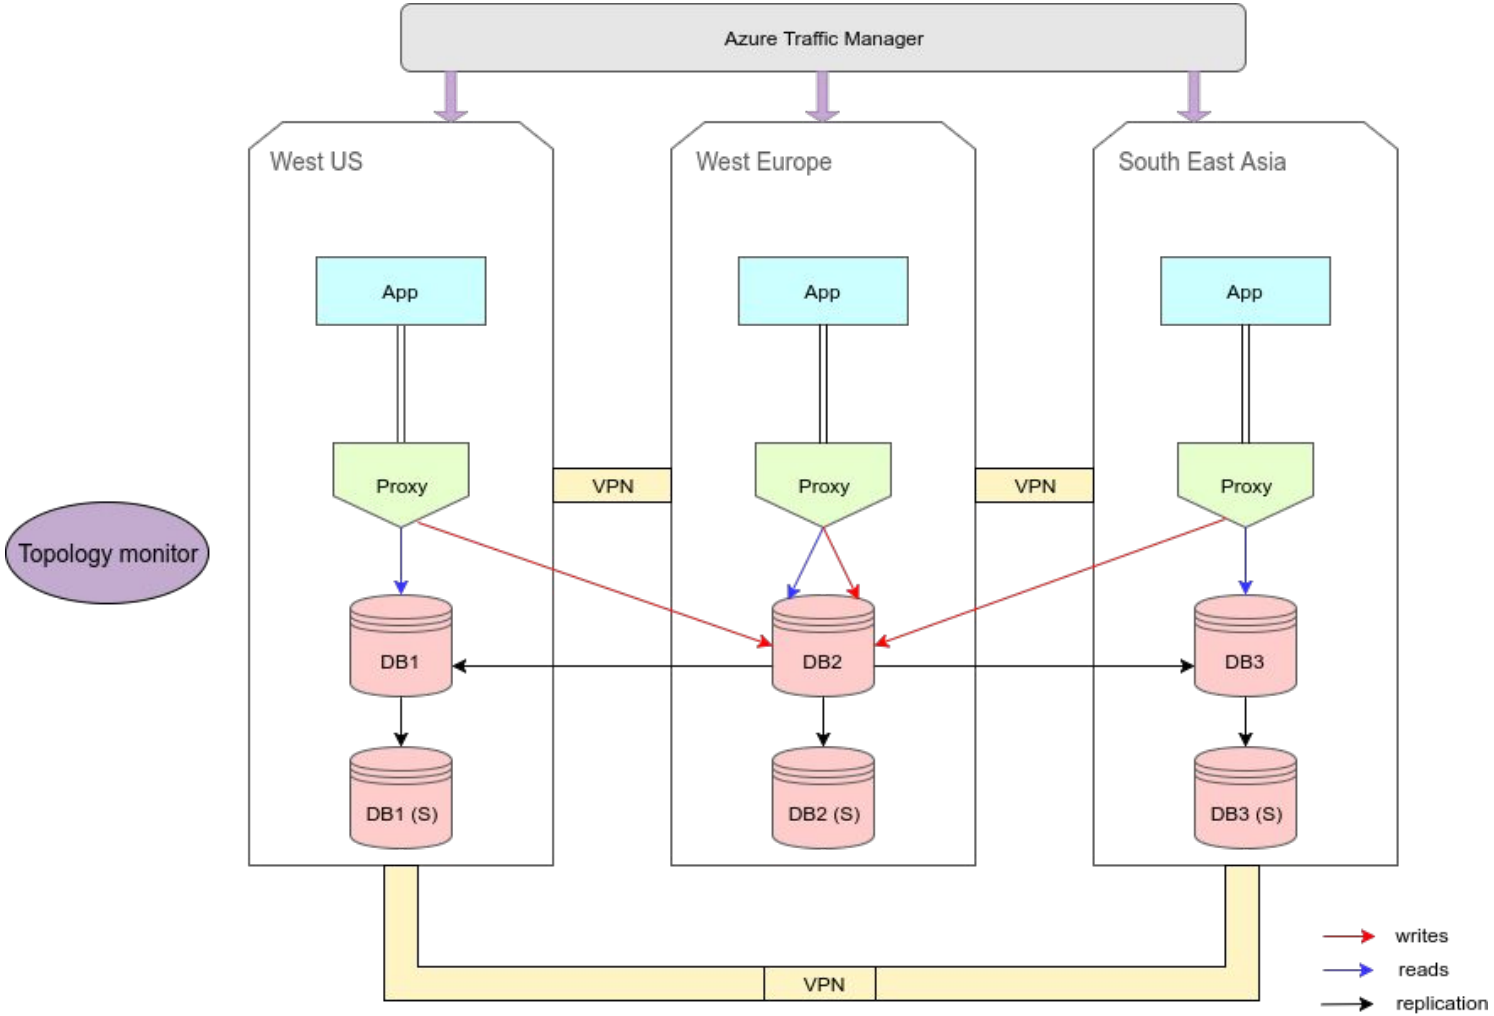

Failure scenarios

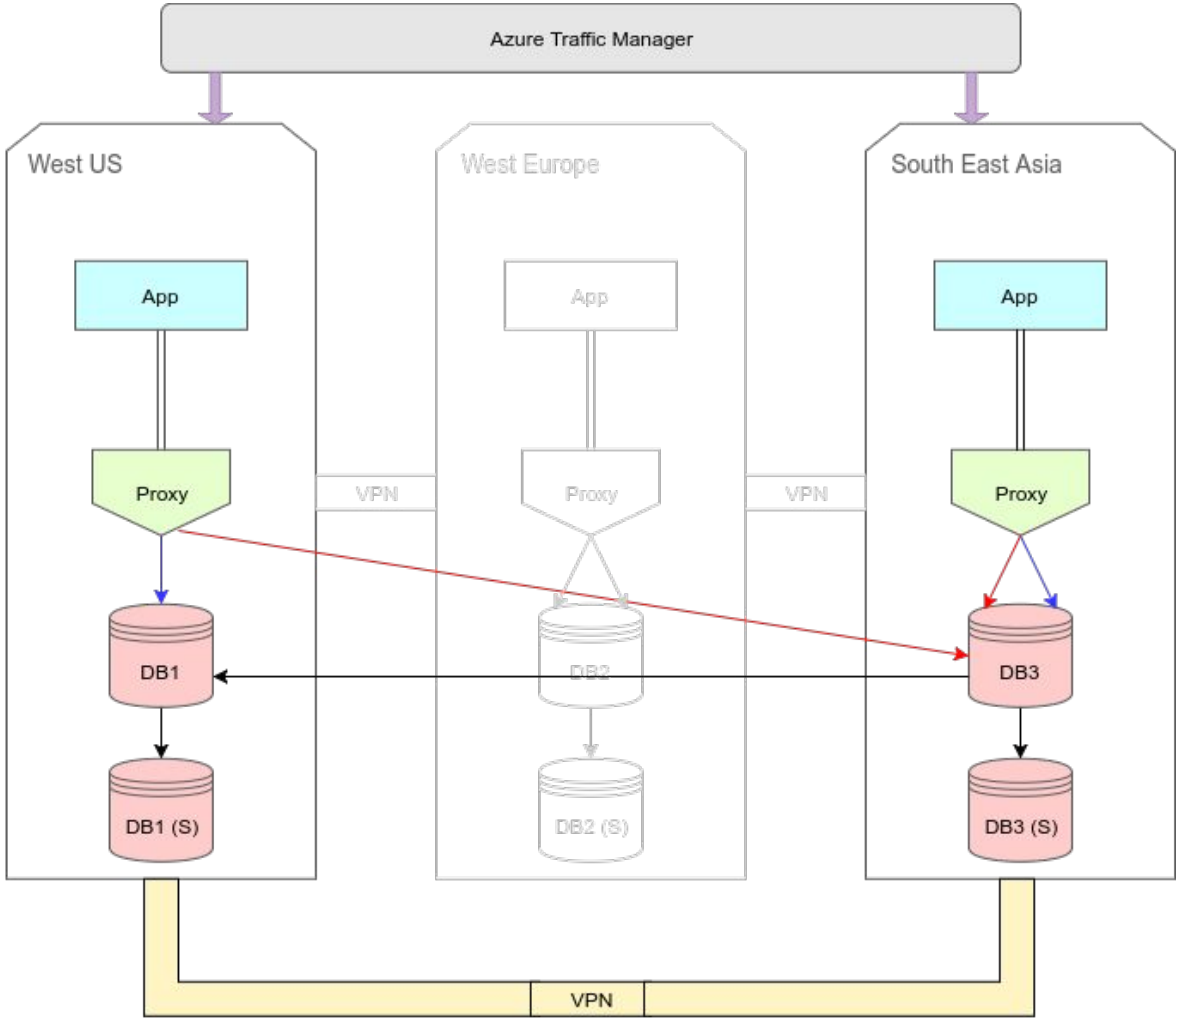

Failure scenarios

- Region failure
  - Traffic manager will redirect traffic to the nearest region (endpoint monitoring)
  - If it is the master region, MHA will promote a new master and the proxies will redirect the traffic accordingly
- Application layer
  - Autoscaling
- Proxy layer failure
  - Traffic manager will redirect traffic to the closest region (endpoint monitoring)

Failure scenarios

- MySQL master failure
  - MHA will promote one of the remaining slaves as the new master
  - ProxySQL will redirect writes to the new master and reads for the affected region
- MySQL slave failure
  - ProxySQL will redirect reads to another available slave

Limitations

Limitations

- No HA at the proxy level
- Writes need to cross a WAN link for two out of three regions
- Read-only based roles
- Fencing could block failover
- MHA manager needs to be restarted after failover

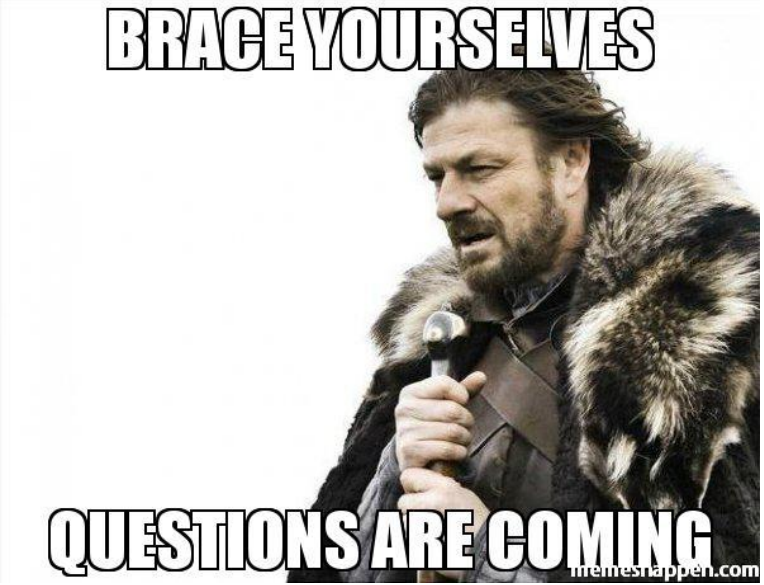
BRACE YOURSELVES
QUESTIONS ARE COMING
memesnappen.com

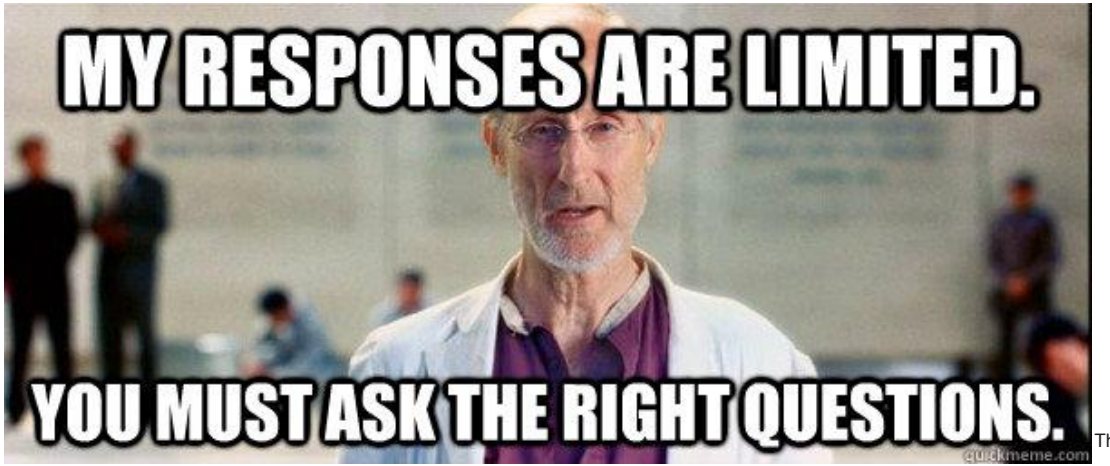
MY RESPONSES ARE LIMITED.
YOU MUST ASK THE RIGHT QUESTIONS.
quickmeme.com

THANK YOU!

Gabriel Ciciliani
Internal Principal Consultant
ciciliani@pythian.com
@gabocic