
Question

[Timur Safin](#) · Jun 15, 2016

Cluster deployment recipes?

Let assume you have a infinitely scaling algorithm implemented in your application, using replication, ECP, or any other means of horizontal scaling, and let assume you know how to run your system under any volume of requests, the trick is to deploy required number of computing nodes in the cluster. If we are talking about cluster of 2-4 nodes your administrator (or as they call it today "devops engineer") will install anything manually. Probably he will easily handle 5 nodes configuration in the cluster. But what if you need to deploy 10, 100, or 200 properly configured, interconnected nodes?

Devops are lazy, if there is some easy way to automate task then they will automate, even if there is need to use newer and fancy tool. [Though requirement here, this new and fancy tool shold be easy to learn: they could spend a couple of hours reading StackOverflow or vendor forums, but certainly have no times for several days. Thus simpler way usually wins]

So, it's easy to replicate nodes in the cluster (be it on-premise or in the cloud) if nodes are identical. vSphere cmdlet, [docker hub images](#), [Digital Ocean droplet](#), or [AWS template](#) all allow to make new, identical instance to be up'n'running in a few moments.

But for ECP I need further fune-tuning during configuration time:

- let assume I've implemented map-reduce controller for which we need to instantiate 1 master node, and N worker nodes.
- Each worker node should mount remotely, via ECP, some particular database from master, where control code and control global are residing.
- In reverse master should mount remotely, via ECP, each worker database, mapping particular global subscript from this database (let's name it "backward data channel").

For 2-3 nodes in the cluster we could make those changes in cache.cpf manually, but for larger scale we should probably use some wiser provisioning tools. Looks like [Ansible](#), as described by Murray before, is a good candidate for such task, and is able to provision images created very flexibly.

But, being lazy and having no experience with Ansible, I'd prefer to see easy to hack recipe. So, have you already cooked anything similar, which will be able to easy configure ECP interconnect while instantiating multiple nodes in the cluster? Anything for vSphere, Digital Ocean, Amazon?

[#Cloud](#) [#ECP](#) [#Caché](#)

Source URL: <https://community.intersystems.com/post/cluster-deployment-recipes>