

Discussion

[Erik Svensson](#) · Sep 18, 2020

## InterSystems Mirroring of code databases

Hello!

First of all, let me state that I am no senior InterSystems expert.

In my organization, we have a HealthShare Health Connect setup where each namespace has one code database and one data database, which are both actively mirrored. We have two nodes in the mirror.

We had a controlled failover last night to make sure that the backup node works as intended, which it didn't. It turned out that we had only deployed code onto the primary node in several namespaces causing errors with missing classes after the failover. So it seems that each time you deploy productions, you have to manually deploy it to both instances (the primary and failover). That makes me wonder:

1. What is actually mirrored when you mirror a code database
  1. Obviously not new classes
  2. changes to existing classes?
  3. settings on the production adapters?
  4. something else?
2. How do you guys go about deploying new code?
  1. Are you utilizing some kind of automation tool to keep the mirrored nodes consistent regarding code and versions of code?
  2. Are you just manually deploying to each node and have good routines doing it?
3. Or do we have some kind of faulty setup which makes this not work as intended?

I don't think our setup is faulty, I think we just missed this a bunch of times which makes me want to abstract this to a way where you deploy to something that deploys the same code to both nodes.

An example: We have 3 environments (production, QA and test), for each of QA and prod, we receive webservice requests from 2 different networks, an internal network and an external one. For each network, we have a pair of web servers running httpd with web gateway. This makes 4 web server hosts for each production and qa environment and in the test environment, we have slimmed this to only have the one pair. Totally 10 web servers. This is bound to be time consuming to maintain and create inconsistency and details that is not done exactly the same between the hosts if you are not extremely thorough, if you would do it manually. So we use Ansible. I have made a playbook and a set of configs for each environment and each network type so each pair is treated exactly the same and the playbook is always used to deploy changes and keep consistency.

I would like to achieve something similar with deploying code to our HeathConnect mirrored instances.

How do you guys do it?

[#Mirroring #HealthShare](#)

---

Source URL: <https://community.intersystems.com/post/intersystems-mirroring-code-databases>