

Help understanding node.id in cluster



[jbradley](#) 60 posts since

Mar 24, 2009

Hi Jive,

We are experiencing some issues in our dev environment. We have 3 node clusters in integration, qa, and staging. It appears the staging instances are trying to join the qa cluster which is already at its license limit (3).

I'd like to understand the node.id (sbs/applications/home/node.id).

Should this value be different for each node in a cluster or the same? Specifically, would you expect 3 different node.ids or a shared node.id for a given cluster?

How is the node.id created? And is it possible to change it? If so, how do we do this?

In our integration environment all three nodes share the same node.id.

In our QA and staging environment, all 6 nodes (3 qa & 3 stg) share the same node.id.

Thanks in advance,

jb



[todd.west](#) 15,926 posts since

Jul 14, 2008 1. **Re: Help understanding node.id in cluster** Sep 18, 2009 7:55 AM

Help understanding node.id in cluster

Hi Justin,

You should not modify the node.id values as these are automatically generated by the application. I believe they are mostly used by us when setting up a Bridged instance such that each node can be easily identified. If you are having issues with getting clustering set up, this is separate from the node.id value.

It sounds like all of your servers reside on the same network/subnet? If so, this is why they are automatically attempting to join together. The multicast discovery of coherence will automatically try to find other nodes to connect to. There are a few ways to get around this:

- 1) Move the servers into different subnets so they are unable to communicate with each other
- 2) Use unicast instead of multicast and specify the server locations directly
- 3) Modify some JVM arguments to allow for the multicast servers to only connect to the servers they should connect to. More information on this here:

<http://www.jivesoftware.com/jivespace/docs/DOC-1940>

Hope that helps. Thanks.

-Todd



[jbradley](#) 60 posts since

Mar 24, 2009 2. **Re: Help understanding node.id in cluster** Sep 20, 2009 5:40 PM

Help understanding node.id in cluster

Thanks Todd, this is great information. My first choice is option 3, modifying JVM arguments. I have two additional questions 😊.

1) Is this the correct file to add the JVM arguments - /usr/local/applications/sbs/bin/instance ?

2) In the above file, I see two interesting variables - APP_CLUSTER_ADDR & APP_CLUSTER_PORT. Do you know if modifying these variables will have the same result as adding tangosol.coherence.localhost & tangosol.coherence.clusterport?

Thanks,

jb



[todd.west](#) 15,926 posts since

Jul 14, 2008 3. **Re: Help understanding node.id in cluster** Sep 21, 2009 6:18 AM

Ah, I didn't realize that you were on the 3.0 platform. That makes it even easier 😊

It's exactly what you expected, all you need to do is set the APP_CLUSTER_ADDR and APP_CLUSTER_PORT and everything should automatically hook up correctly. Let me know if you have any issues with this. Thanks.

-Todd



[jbradley](#) 60 posts since

Mar 24, 2009 4. **Re: Help understanding node.id in cluster** Sep 21, 2009 2:46 PM

Hi Todd,

Help understanding node.id in cluster

I think we have it working. All nodes are clustering on the correct dev environment. I only changed the APP_CLUSTER_PORT argument, and left the APP_CLUSTER_ADDR argument as is.

So all instances are using the default address: 224.224.224.224

And each dev environemnt is assigned a different port: 9003 (stg), 9004 (int), or 9005 (qa).

I'll keep my eye on the logs to make sure all are happy, but does this sound like it should work to you?

Thanks,

jb



[todd.west](#) 15,926 posts since

Jul 14, 2008 5. **Re: Help understanding node.id in cluster** Sep 21, 2009 2:48 PM

Yep, that should work out perfectly! Let me know if you run into any issues and I'd be happy to help. Thanks.

-Todd