

Re: Meanings of thresholds in the cluster administration snap-in

Source:

<http://www.tech-archive.net/Archive/Windows/microsoft.public.windows.server.clustering/2008-07/msg00146.html>

- *From:* setecastronomy <setecastronomy@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Thu, 24 Jul 2008 08:56:01 -0700
-

Please, can you say something more about the "RetryPeriodOnFailure".
How vcan I change enable it in 2003 EE ? Is it a registry key ?

As far as I know, there is no way to delay these restart attempts in cluster.

I asked about it and threshold because I'm plagued by a problem with a service. In our server architecture under particular circumstances it cannot go online for few minutes for acceptable reasons. Unfortunately the cluster tries to restart it at the speed of light (or at least it is what I deduced from watching event viewer) so the threshold is reached no matter how high I can set it. If only I could delay these restart attempts I could avoid the few minutes backout of our architecture.

By the way the problem is originated by SQL server running under the domain account 'cluster' so if the dc cannot authenticate him the service cannot go online. The unavailability of DC is only temporarily but the cluster is faster to reach that damn threshold. After that the cluster may be healthy but it will not try to restart or failover longer.

Thanks
Filippo

"John Toner [MVP]" wrote:

If your SQL server resource fails to start on a node, MSCS will immediately attempt to restart that resource x number of times until it hits the threshold values. Once the threshold value is reached, the resource will either cause the whole group to move to another node (default behavior) or it will fail if all thresholds are hit. There is a Group level threshold also in place so that a failing resources in a cluster does not cause the group to "ping pong" eternally between the nodes.

Once all thresholds are hit (resource and group level thresholds) the resource will fail and remain in a failed state until you try to restart the

Re: Meanings of thresholds in the cluster administration snap-in

resource again. You can have cluster periodically attempt to restart the resource by setting a "RetryPeriodOnFailure" value...by default, this is disabled in 2003 (set to 60 minutes in 2008).

As far as I know, there is no way to delay these restart attempts in cluster. Service control manager settings do not apply to the cluster...it will attempt to restart immediately.

I hope this helps.

Regards,
John

Visit my blog: <http://msmvps.com/blogs/jtoner>

"setecastronomy" <setecastronomy@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message news:A8EB6359-F6D8-4191-B422-B35680A4FEE8@xxxxxxxxxxxxxxxxxxxx

I have SQL server 2000 running on cluster realized using 2003 EE. The

cluster

is active/passive, SQL is installed on both nodes while its data file

resides

on a NAS.

It works.

Using cluster administration snap-in you have the possibility, for each group of resources, to set a threshold and a period for failover.

As far as I know it means the cluster will try to bring online the group

on

the other node only if the the number of failures is less than the threshold. If, for some reasons, it reaches the threshold, it gives up and doesn't try any more. It may be a smart choice because it is useless to

try

indefinitely to bring online resources that fail permanently due to

serious

reasons, with the risk to compromise and overload the entire cluster.

For each resource inside a group, in the advance tab, there is the possibility to set "Restart" behaviour with an associated threshold and period.

I think the meaning is that the node will try to restart the failed

Re: Meanings of thresholds in the cluster administration snap-in

resource

till it reaches the threshold.

What it is not clear to me is how these two threshold are related and when they cause the switch on the other node.

Focusing on my SQL server, what happens if it fails to start on a node ? I suppose the node try to restart it till it reaches the threshold fixed for that resource and only then a failover is initiated on that group. Is it correct ?

The most important point I'm concerned with is that I would like to fix

the

interval of time between a failure and a restart attempt. Windows let me configure it using the service.msc but I doubt it is the right way to do operate on clusters, where you should use the cluster administration

snap-in.

I searched on Microsoft's site a lot but I found only a brief description

of

all these parameters and no one complex scenario where they all interact. Most documents only say that in case of failure the resource is moved to

the

other node, but how this fact is influenced by threshold and restart

policies

is not analyzed in depth.

Can you give me a hand ?