Re: Long failover time...

Source:
http://www.tech−archive.net/Archive/Windows/microsoft.public.windows.server.clustering/2007−08/msg00117.html

- *From*: "Rodney R. Fournier [MVP]" <rod@xxxxxxxxxxxxxxxxxxxxxxxxxxx>
- *Date*: Mon, 13 Aug 2007 15:22:40 −0500

Yeah I missed the 90% one. But I will still contain that it (hosts file) does nothing with a properly configured systems and DNS. I think you need to check into bindings, cause I believe your are not setup correctly.

As I stated both nic will use the hosts files and each will need the address of the other one. In your example you hard code the heartbeat? I would love to see your nic configurations, bindings, etc. because you have to be missing a best practice or 3 to have to use this kind of host file! I hope nobody tries this that reads your post!

MSCS is only Highly Available. Parts of any system, Exchange, SQL can be FT − Raid hardware, redundant power supplies, etc. The complete system will only be HA.

Cheers,

Rodney R. Fournier

MVP − Windows Server − Clustering
http://www.nw−america.com − Clustering Website
http://msmvps.com/clustering − Blog
http://www.clusterhelp.com − Cluster Training
ClusterHelp.com is a Microsoft Certified Gold Partner

"Kenny Speer" <kenny.speer@xxxxxxxxx> wrote in message
news:OolSpPe3HHA.1204@xxxxxxxxxxxxxxxxxxxxxx

> Ah, first it's not 90% of clusters, it's 90% of clusters with long takeover times. Your context is wrong.
>
> Second, in your hosts file, you put a hostname and IP address. Each NIC only has one IP address (in this config, yes it's possible to add more), therefore, the hosts file is used for that network only. For example:
>
> Public Network: 10.x.x.x
> Private Cluster Network: 192.168.x.x

In your hosts file, you add: Node1 192.168.x.1 and Node2 192.168.x.2

Pretty simple. If Node1 wants to communicate with Node2 for any reason, it goes over the private cluster network. No DNS required. This will ensure that your cluster is running regardless of your DNS status. For any other network comm, normal DNS lookups are used.

You distinguish between Highly Available and Fault Tolerant (yadda yadda) here is a Microsoft page which uses Fault Tolerant quite often:

http://technet.microsoft.com/en−us/library/aa997507.aspx

http://technet.microsoft.com/en−us/library/aa997234.aspx

As for where I read it, /etc/hosts and drivers/etc/hosts have been used forever to bypass DNS in mission critical scenarios.

"Rodney R. Fournier [MVP]" <rod@xxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message news:e9oQs%23d3HHA.4400@xxxxxxxxxxxxxxxxxxxxxx

> Kenny, I am not about to argue that DNS can be an issue. I am firmly stating that I have never needed a host file on a clustered node. Does not matter who runs DNS, they just need to ensure it works. DNS is really easy stuff to get right or wrong. If you really have had a need on 90% of clusters for DNS, then they had DNS issues, plain and simple. It is not a best practice to create a hosts file for a clustered node. Does not matter if DNS is Microsoft or not, it needs to be configured correctly.
>
> Next, you mentioned a hosts file for the private only? Where did you read about this? How do you tell a hosts file to only work for one NIC and not the other(s)?
>
> MSCS is not and never will be fault tolerant, it is considered highly available. And Windows 2003 with or without clustering is dependant on DNS, Active Directory has for a long time now.
>
> Cheers,
>
> Rodney R. Fournier
>
> MVP − Windows Server − Clustering
> http://www.nw−america.com − Clustering Website
> http://msmvps.com/clustering − Blog
> http://www.clusterhelp.com − Cluster Training
> ClusterHelp.com is a Microsoft Certified Gold Partner
>
> "Kenny Speer" <kenny.speer@xxxxxxxxx> wrote in message news:eJwgbid3HHA.5852@xxxxxxxxxxxxxxxxxxxxxx

Not true. If the DNS server is running properly you won't have an
issue. A perfectly configured Cluster can have any number of issues due
to a misconfigured/non functioning DNS server. To say that 90% of
clusters were misconfigured is complete crap. Often, an admin who owns
the clusters does not own the DNS servers and often those DNS servers
are not Microsoft servers. It is absolutely an option to put the
cluster host names ONLY in the hosts file. Scalability is not an issue,
how many nodes are in your cluster? The majority of clusters are 2−4
nodes, 4 lines isn't very many to add to a host file which should not
change.

Also, remember, the hosts file has the private interface addresses ONLY,
not the public. This insures that all cluster comm from one node to the
other will use the cluster comm interface (i'm not just talking about
heartbeat here) without requiring or *depending* on an outside service
(DNS).

Anyway, I think your argument doesn't hold any water, since we're
talking about making MSCS as fault tolerant as possible but then you
make it *depend* on DNS for proper operation (not access from clients,
but even then, we've used IP addresses for 30 years when DNS goes down).

~kenny

"Rodney R. Fournier [MVP]"
<rod@xxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
news:eDsCWQd3HHA.4436@xxxxxxxxxxxxxxxxxxxxxxx

> Than 90% of clusters were configured
> wrong to begin with :) Honestly,
> if DNS is running properly you won't ever
> need a hosts file. Host files
> don't scale very well or allow for easy
> changes!

Cheers,

Rodney R. Fournier

MVP − Windows Server − Clustering
http://www.nw−america.com − Clustering Website
http://msmvps.com/clustering − Blog
http://www.clusterhelp.com − Cluster Training
ClusterHelp.com is a Microsoft Certified Gold Partner

"Kenny Speer" <kenny.speer@xxxxxxxxx>
wrote in message
news:OvjOjNd3HHA.1208@xxxxxxxxxxxxxxxxxxxxxx

> Really? Hmmm. I've seen
> 90% of long failover times
> resolved by doing
> this.
>
> ~kenny

> "Rodney R. Fournier
> [MVP]"
> <rod@xxxxxxxxxxxxxxxxxxxxxxxxx>
> wrote in
> message
> news:OdC0XMc3HHA.5984@xxxxxxxxxxxxxxxxxxxxx
>
> > Hosts file!
> > No, DNS
> > works
> > nicely :)
> >
> > Cheers,
> >
> > Rodney R.
> > Fournier
> >
> > MVP −
> > Windows
> > Server −
> > Clustering
> > http://www.nw−america.com

− Clustering
Website
http://msmvps.com/clustering
− Blog
http://www.clusterhelp.com
− Cluster
Training
ClusterHelp.com
is a
Microsoft
Certified
Gold
Partner

"Kenny
Speer"
<kenny.speer@xxxxxxxxx>
wrote in
message
news:uaFNDsH2HHA.5884@xxxxxxxxxxxxxxxxxxxxxxx

I
thought
it
was
recommended
to
put
all
cluster
names/ips
in
the
hosts
file
located
at:
%WINDIR%\system32\drivers\etc\hosts

By
default,
you
don't
want
your
cluster
communicating
over
the
public

(client
access)
interface
and
you
don't
want
your
cluster
to
fail
just
because
your
DNS
server
goes
down
or
is
not
accessible.

By
adding
each
node
to
this
file
on
both
nodes,
you
won't
even
do
a
name
lookup
via
DNS
since
Windows
uses
hosts/lmhosts
then
DNS.

~kenny

"LOVEBEINGDBA"
<LOVEBEINGDBA@xxxxxxxxxxxxxxxxxxxxxxxxx>
wrote
in
message
news:EC8A3B7D−E8D7−4B1E−BF4A−DD63E31750E9@x

Thanks
and
I
appretiate
your
reply...

Is
it
required
to
create
a
PTR
record
for
the
Cluster
Name
in
the
DNS
Server???
We
already
have
a
cluster
whose
name
is
registered
in
the
DNS
only
for
forward
lookup...

Thanks.
Arun
M

"John Fullbright" wrote:

> "The server for 143.3.16.172.in−addr.arpa. could not be contacted over adapter 'Public' to determine whether it accepts DNS registration updates. Retrying at a later time"
>
> Looks like problems contacting an authoratative DNS server for the reverse lookup zone that holds the PTR record.
>
> ..

"LOVEBEINGDBA"
<LOVEBEINGDBA@xxxxxxxxxxxxxxxxx
wrote
in
message
news:A70ABBFB−6DDA−495D−B715−542

Environment:
Windows
Server
2003
R2
x64
SP2
MSCS
2
node
failover
cluster

Failover
takes
about
2
minutes.
When
doing
the
failover
the
clustername
take
a
long
time
to
come
back
up.

Clustername
has
been
registered
in
the
DNS
for
forward
look
up

.

We have another production cluster and that works fine in the same setup.

These are the lines I found that are related to this issue from cluster log:

00000840.00000bb4::2007/08/06−16
WARN Network Name <Cluster Name>: The server for 143.3.16.172.in−addr.arpa. could not be contacted over adapter 'Public' to determine

whether
it
accepts
DNS
registration
updates.
Retrying
at
a
later
time.
00000840.00000bb4::2007/08/06−16
INFO
Network
Name
<Cluster
Name>:
Replaced
DNS
name
<clustername.domain>
with
IP
Address
172.16.3.143
over
adapter
'Public'.
00000840.00000b08::2007/08/06−16
INFO
Network
Name:
time
until
next
DNS
reg:
2007/08/06−22:17:35
(128309122556769715)
00000840.00000bb4::2007/08/06−16
WARN
Network
Name
<Cluster
Name>:
Failed
to
register
DNS
PTR
record

143.3.16.172.in−addr.arpa.
for
host
<clustername.domain>
over
adapter
'Public',
status
1460
00000840.00000bb4::2007/08/06−16
INFO
Network
Name
<Cluster
Name>:
Modified
DNS
name
<clustername.domain>
with
IP
Address
172.16.3.143
over
adapter
'Public'.
00000840.00000b08::2007/08/06−16
INFO
Network
Name:
time
until
next
DNS
reg:
2007/08/06−22:17:35
(128309122556769715)

Any
help
will
be
greatly
appretiated.

.