

Re: Long failover time...

Source:

<http://www.tech-archive.net/Archive/Windows/microsoft.public.windows.server.clustering/2007-08/msg00123.html>

- *From:* "Kenny Speer" <kenny.speer@xxxxxxxx>
 - *Date:* Mon, 13 Aug 2007 17:37:40 -0400
-

So we resort to insults? You've no idea what my background is or how long I've been working with TCP/IP and your comment has nothing to do with the topic of discussion.

Very nice. I don't think Comer is going to discuss heartbeat messages.

Again, you guys are way over reacting. CCS uses the host file, exchange (edge nodes) use the host file, etc etc.

And, I was explicitly stating that the cluster comm was NOT used just for the heartbeat. It's used for all cluster communication. So, you're right, I'm not arguing, heartbeat doesn't care about the hostname, and I never said it did. To refresh you both "(i'm not just talking about heartbeat here)"

"Rodney R. Fournier [MVP]" <rod@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message <news:uktgK2e3HHA.4676@xxxxxxxxxxxxxxxxxxxxxxxx>

Kenny, I suggest you read Comer's book TCP/IP. Hosts files are not needed for the heartbeat and Russ and I have already stated. If you can state any Microsoft white papers on the subject, I am all ears!

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@xxxxxxxx> wrote in message <news:OVhAvse3HHA.5880@xxxxxxxxxxxxxxxxxxxxxxxx>

Well, I've been setting the IPs in the host file for 7 years now and have not run into one issue. Again, you only set the IPs of the private interfaces and, AFAIK, all best practices show the private interfaces as not being routed anyway, so there should be no issue whatsoever.

Re: Long failover time...

I think you're arguing on documentation and best practices which don't always consider the optimal solution. They consider the *best* solution for the majority of users based on simplicity and stability.

As for the hosts file. A NIC doesn't use the host file. The routing stack in the host uses it in the following order: lmhosts, hosts, dns and wins. There is no issue with bindings or nic configurations. This is a very simple task that effects nothing except for the IP for the hosts you put in the file. So, the normal lookup for an IP addr is like this:

Node1 -> where is Node2 -> is it cached? nope-> how do I find it? -> go to hosts/lmhosts -> nothing there (they are empty by default) -> use dns/wins -> found it, use IP returned by gethostbyname() ... The only difference now, is it finds the host IP in the hosts file and no network requests are required.

Again, this does not effect anything except for how an IP address is returned to from the function gethostbyname(), doesn't effect routing, nic config, etc etc. And it only effects those hosts you put in the file, hence you only add the cluster Nodes in the file.

~kenny

"Rodney R. Fournier [MVP]" <rod@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>
wrote in message news:uOpezee3HHA.5796@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx

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

Re: Long failover time...

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

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@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>
wrote in message
news:e9oQs%23d3HHA.4400@xxxxxxxxxxxxxxxxxxxxxxxxxxxx

Kenny, I am not about to argue that DNS can be an

Re: Long failover time...

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

Re: Long failover time...

"Kenny Speer"

<kenny.speer@xxxxxxxx>

wrote in message

news:eJwgbid3HHA.5852@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx

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

Re: Long failover time...

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

Re: Long failover time...

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@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>

wrote in
message

news:eDsCWQd3HHA.4436@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

Than
90%
of
clusters
were
configured
wrong
to
begin
with
:)
Honestly,
if

Re: Long failover time...

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

Re: Long failover time...

Gold
Partner

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

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

~kenny

"Rodney
R.
Fournier
[MVP]"
<rod@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>
wrote
in
message
news:OdC0XMc3HHA.5984@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

Hosts
file!
No,
DNS
works
nicely
:)

Cheers,

Re: Long failover time...

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@xxxxxxxx>

wrote
in
message
<news:uaFNDsH2HHA.5884@xxxxxxxxxxxx>

I
thought
it
was
recommended
to
put
all
cluster
names/ips
in
the
hosts

Re: Long failover time...

file
located
at:
%WINDIR%\system32\drivers\etc\h

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

Re: Long failover time...

even
do
a
name
lookup
via
DNS
since
Windows
uses
hosts/Imhosts
then
DNS.

~kenny

"LOVEBEINGDBA"
<LOVEBEINGDBA@XXXXXXXXXXXXX
wrote
in
message
news:EC8A3B7D-E8D7-4B1E-BF

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

Re: Long failover time...

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
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

Re: Long failover time...

authoritative
DNS
server
for
the
reverse
lookup
zone
that
holds
the
PTR
record.

..

"LOVEBEINGDBA
<LOVEBEINGDBA

wrote
in
message
news:A70ABBFB-6

>
Environment:

>
Windows
Server
2003
R2
x64
SP2

>
MSCS
2
node
failover
cluster

>
>
Failover
takes
about
2
minutes.

When
doing
the
failover
the
>
clustername

Re: Long failover time...

>
> 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

Re: Long failover time...

Re: Long failover time...

related
to
this
issue
>
from
cluster
>
log:
>
>
00000840.00000bb4
WARN
Network
Name
>
<Cluster
>
Name>:
>
The
server
for
143.3.16.172.in-add
could
not
be
>
contacted
over
>
adapter
>
'Public'
to
determine
whether
it
accepts
DNS
registration
>
updates.
>
Retrying
>
at
a
later
time.
>

Re: Long failover time...

Re: Long failover time...

00000840.00000bb4
INFO
Network
Name
>
<Cluster
>
Name>:
>
Replaced
DNS
name
<clustername.domain
with
IP
Address
>
172.16.3.143
over
>
adapter
'Public'.
>
00000840.00000b08
INFO
Network
Name:
>
time
until
>
next
DNS
reg:
2007/08/06-22:17:3
(1283091225567697
>
00000840.00000bb4
WARN
Network
Name
>
<Cluster
>
Name>:
>
Failed
to
register
DNS
PTR

Re: Long failover time...

Re: Long failover time...

```
record
143.3.16.172.in-addr
>
for
host
>
<clustername.domain>
over
adapter
'Public',
status
1460
>
00000840.00000bb4
INFO
Network
Name
>
<Cluster
Name>:
>
Modified
DNS
name
<clustername.domain>
with
IP
Address
>
172.16.3.143
over
>
adapter
'Public'.
>
00000840.00000b08
INFO
Network
Name:
>
time
until
>
next
DNS
reg:
2007/08/06-22:17:30
(1283091225567697)
>
>
```

Re: Long failover time...

Re: Long failover time...

Any
help
will
be
greatly
appretiated.