Re: Best replication architecture?

Source:
http://www.tech−archive.net/Archive/SQL−Server/microsoft.public.sqlserver.replication/2006−08/msg00037.html

- *From*: Lauren <Lauren@xxxxxxxxxxxxxxxxxxxxxxxxx>
- *Date*: Wed, 19 Jul 2006 06:51:01 −0700

I have just been reading about alternate sync partners which is not something I realised would be needed. From what I can see it seems to have been deprecated for SQL Server 2005 without any alternative having been provided to do the same job. I definitely will need my big Express clients to be able to subscribe to any other server in their domain if the link goes down for some reason.

"Hilary Cotter" wrote:

> Yes, that is what I was referring to. You can have a client subscribe to a different publisher. So if it is subscribing to Publisher 1, and the link to Publisher 1 goes down you would have to configure Publisher 2 in advance as an alternate sync partner for Publisher 1, and then manually configure this subscriber to use Publisher 2.
>
> There are some problems with this. The big one is that the Subscriber must reconnect with Publisher 1 before the retention period has elapsed, and must connect back to Publisher 1 to have its identity range incremented. You will have to manually fail it back to Publisher 1.
>
> −−
> Hilary Cotter
> Director of Text Mining and Database Strategy
> RelevantNOISE.Com − Dedicated to mining blogs for business intelligence.
>
> This posting is my own and doesn't necessarily represent RelevantNoise's positions, strategies or opinions.
>
> Looking for a SQL Server replication book?
> http://www.nwsu.com/0974973602.html
>
> Looking for a FAQ on Indexing Services/SQL FTS
> http://www.indexserverfaq.com

"Lauren" <Lauren@xxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message news:7DE3EE64−318C−4F58−9B50−FB06B6AF4073@xxxxxxxxxxxxxxx

This is very helpful, Thankyou

I have more questions though ...

Do you mean a Master server for each domain (which may have one or more installations within it)?
In our scenario, connections could be very unstable at times so we would need the clients to be able to connect to different servers if their nominated one is unavailable.
Also it is only the 4 GB client machines which require a subset of data, the
servers and large Express machines should be identical to one another.

"Hilary Cotter" wrote:

> You are correct but it sounds like you are using MS Exchange terminology.
>
> Basically you will have a "master" server or a top level publisher
> which
> is similar to a MS Exchange bridgehead server. This server will publisher
> to
> the remaining 5 SQL Servers which each have their own domain (Publisher1,
> Publisher2, Publisher3, Publisher4, Publisher5). The filter should be
> constructed so that only data for publisher 1 goes to publisher 1 and
> publisher 2 goes to publisher 2. If your don't filter enough you won't be
> using your LAN efficiently. It can work mind you, but then data will be
> going to both Publisher 1 and 2 to the main publisher and back and forth.
> It
> is acceptable to have this overlap, just not terribly efficient.
>
> Publisher 1 will then publish to a group of SQL Server Express clients.
> Publisher 2 will publish to another group.
>
> Publisher 1 SQL Server Express subscribers should always replicate with
> Publisher 1 and not Publisher 2. Again you can have some overlap, but its

not very efficient. Unless you are under significant load you should not
have a problem.

You can to a degree have Publisher 2 back up Publisher1 but its not
optimal.

Also the master server would not be that heavily loaded in a hierarchy.
It
would be more heavily loaded if you did not have a hierarchy.

––
Hilary Cotter
Director of Text Mining and Database Strategy
RelevantNOISE.Com − Dedicated to mining blogs for business intelligence.

This posting is my own and doesn't necessarily represent RelevantNoise's
positions, strategies or opinions.

Looking for a SQL Server replication book?
http://www.nwsu.com/0974973602.html

Looking for a FAQ on Indexing Services/SQL FTS
http://www.indexserverfaq.com

"Lauren" <Lauren@xxxxxxxxxxxxxxxxxxxxxxxxx> wrote
in message
news:8060857D−6740−4E40−A4D9−2B7A62AD9B9E@xxxxxxxxxxxxxxx

> Thanks for your reply Hilary.
>
> Each server in our system is an installation.
> The system we are providing will have these
> installations in different
> physical locations at different times.
> Sometimes it is envisaged that
> each
> installation will be stand alone in its own
> domain and at others the
> user
> will be able to link one or more installations
> to each other within a
> single
> domain. The design proposed in the project

plan is that all servers
within
a
domain (where there is more than one)
would replicate to each other and
a
bridgehead server would be nominated that
would replicate to the other
domains. This sounds similair to what you
have suggested but your
suggestion
sounds more efficient than having all−to−all
replication going on
within a
domain. Would the master server be very
heavily loaded or could that be
solved by using pull subscriptions?

Sorry if my questions are naive but I am still
in the very early stages
of
trying to read and understand this problem.

Thankyou,
Lauren

"Hilary Cotter" wrote:

> Answers inline.
>
> −−
> Hilary Cotter
> Looking for a SQL Server
> replication book?
> http://www.nwsu.com/0974973602.html
>
> Looking for a FAQ on
> Indexing Services/SQL FTS
> http://www.indexserverfaq.com
>
> "Lauren"
> <Lauren@xxxxxxxxxxxxxxxxxxxxxxxxx>
> wrote in message
> news:D00BC0D8−6990−4968−93F2−F9502508A9F6@xxxxxxxxxxxxxxx
>
> > Hi,
> >
> > I have just
> > started on a
> > new project

and have the responsibility of setting up the SQL Server replication architecture. I wonder if anyone can give me advice on the way to go? We have a setup of 6 SQL Server 2005 servers which should all be replicating to each other, connected to these will be a number of smaller groups of machines (1 "server" and 9 "clients") which will all be running SQL Server Express. The Express "server" has at least 80Gb and a reasonably powerful processor. The Express "clients"

only have a
4Gb Hard
drive which
is
a
bit
of
a limitation
to say the
least but
non−negotiable.
The Express
machines
may
go
offline
deliberately
but the user
still needs
to be able to
work
when
this
happens and
so conflicts
will need to
be resolved
on
re−connect.
The
6
SQL
Server
machines
all need to
be up to
date so that
the Express
"server"
can
connect to
any one that
is available
to it and see
the latets
information.
Oh, and the
tiny Express
"clients"
will only
require a

subset of the overall data applicable to their user domain. How best to achieve this?

The designer on the project has been inclined towards peer−to−peer transactional replication, but as this does not handle conflict resolution he has tried to impose a nauseating check in/out process (and the concept of a master server for data) on the user in order to prevent concurrent updates. This is not the way we want to proceed now that real development is

taking
place and
we have
decided to
go with an
optimistic
locking
mechanism
with
replication
handling
conflicts
(latest is
best or user
role priority
etc.).

You need to do merge
replication between the
express clients and the
main
SQL Servers. Use filtering,
possibly join filtering or
dynamic join
filtering so the express
clients only get what they
need.

Queued replication will log
the conflicts but you will be
unable to
roll
them back.

I further think, although it is
not competely clear to me
that you
want
to
set up a merge replication
heirarchy between your
main sql servers, ie
one
will be the master, all others
subscribers to the master.
Then you
will
each
of these SQL Servers will
publish to the SQL Express

clients.

Can anyone offer me advice on this − I have been reading a lot about replication (which is not something I have ever implemented before) and overall am leaning towards transactional with queuing and updateable subscriptions but I don't know if Express supports subscriber updateable or not and there are lots of features of merge replication which seem to fit our architecture? I am currently trying to get some machines set up to try

things
out for real
but am
having
difficulty as
our IT
provider has
clamped
down
our
machines so
much that I
can't even
get Express
to install!

TIA,

Lauren

.