

Re: IS replication REALLY the way to go?

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.replication/2006-09/msg00587.html>

- *From:* dbermingham@xxxxxxxxxx
 - *Date:* 28 Sep 2006 13:25:05 -0700
-

We support both Asynchronous or Synchronous mirroring. With synchronous mirroring the data is written to the target first, so the source will have to wait for the target write to complete before the write is committed on the source. This guarantees no data loss. Obviously there is a performance penalty for synchronous replication and it is not recommended across WANs. Even storage based synchronous replication will have some pretty hefty requirements when it comes to the type of connection they require and the latency it can tolerate.

The alternative is Asynchronous replication, which takes a copy of the write and sends it across the wire (keeping write order integrity), however it lets the original write complete without delay. This is the preferred method of replication across a WAN, although in the event of an unexpected failure, you could lose whatever is in queue, which is generally at most the last write. The benefit is that the source server does not receive a performance penalty and this option works just fine over a T1 with high latency.

Failover to the target site and failback for SQL server is automatic and generally in the 30-45 second range, and this includes not only data recovery, but starting SQL, recovering an IP address or doing DNS update, and starting any third party app as well. The mirroring is automatically reversed once the original host comes back online and just the data that has changed since it was offline is then sent back to the original server. Once the mirror is back in a mirroring state, it is as simple as clicking on a button in the GUI to bring the original primary back online.

When you say that with Hitachi or SRDF that the DR site is up immediately, you mean the data, not the applications, correct? Storage based replication is only concerned about the data, not the applications. Once the data is recovered, you are still left with the job of recovering the data. One of my customers was trying to do DR for Exchange with storage based replication, and was having a very hard time with it (<http://www.itreseller.com/pr/6434>). When he learned about SteelEye LifeKeeper, he didn't believe that something so inexpensive would work so much better than the solution he spent 100's of thousands of dollars on.

Re: IS replication REALLY the way to go?

SteelEye Data Replication runs as a Windows filter driver, below the file system. Being at that layer means we take very little system resources, maybe 1% CPU and minimal RAM. Disk activity is only increased during the initial sync when the data needs to be read from disk. After a mirroring state is reached, the disk on the source is not impacted. The only time you will see any additional load on the system is if you turn on data compression. Even at the highest level of compression, on a single processor pentium 4 system you will see at most a 15% CPU utilization increase. Compression is only required if your rate of change exceeds the capacity of the WAN link. A T1 can handle a rate of change of about 650 MB per hour with no compression. Beyond that, you will need compression.

Hilary Cotter wrote:

It depends on your ability to tolerate data loss. Do you have a no data loss guarantee? Hitachi does, you do have to pay for it though.

Another point to keep in mind, with Hitachi or SRDF, if the source goes down the destination/dr site is up immediately with no data loss. Does your product provide such functionality? Failback is similarly uncomplex.

Your product is hostbased, it consumes resources on the host. SRDF, Hitachi and other hardware mirroring systems are not host based.

--

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>

"daveberm" <david.bermingham@xxxxxxxxxxxx> wrote in message
<news:1159452244.134449.223610@xx>

Data replication does not need to cost millions. Host based replication such as SteelEye Data Replication (SDR) and LifeKeeper HA can be had for <\$10,000.

The maximum throughput of a 2Mbps line is 900 MB/hour minus Tcp overhead. It is less if the protocol being used is not optimized for

Re: IS replication REALLY the way to go?

WAN. SteelEye is optimized for WAN and also includes data compression, so depending on how much your data can be compressed (~4X), I would assume you might get about 3.5 GB per hour across the pipe. So the initial sync would take 2+ days. After that, you would only be sending the changes.

If you are based in the UK and would like more info, I suggest you contact Open Minds at <http://www.openminds.co.uk/>

David Bermingham, MCSE, MCSA:Messaging
Senior Systems Engineer
www.steeleye.com

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>

"Methodology"
<Methodology@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in
message
news:397E5FBE-91BA-4791-8289-448D0EFD1C00@xxxxxxxxxxxxxxxxxxxx

so

Im coming to end of my replication testing
and im really wondering if I
want
to do this. some basic problems that im
noticing:

if i create a new table / SP I have to re-run
the snapshot agent to
create
and apply a new snapshot to the subscriber.
But my main dabse is 200GB,
and
my subscriber is on the end of a 2Mb sdsl in
Paris, so unless im being
completely stupid, its gonna take
me..like...17 weeks to get the new
snap
over, correct?

Re: IS replication REALLY the way to go?

If you add an individual article through the stored procedures (as opposed to through the wizard) the snapshot will only be for the new article and related metadata. It shouldn't take 17 weeks – I take it you are exaggerating. You could zip it up send it to the subscriber and then pull the subscription and point to an alternate snapshot download folder.

I COULD create a publication per table, so that when I create a new table, I just create a new pub for it, but doesnt each pub have its own agent? so in going to need a CRAY computer to handle all the overhead, correct?

No, by default all publications will replicate to the same subscription database through a single shared distribution agent.

I admit that I can create a separate pub for my tables and my SP's so that I dont have to snap / rep everything again when I want to create a new SP, but surely theres an easier way than this to do things?

I reindex on a Sunday – what the hell is that going to do with my replication?

It puts added load on the log reader agent. YOu should be fine though.

Re: IS replication REALLY the way to go?

Re: IS replication REALLY the way to go?

Im creating a disaster recovery site for my
London office in Paris. We
use
big fat hairy SQL databases that need an
exact copy running in France –
is
replication REALLY what I need? I
haven't looked at mirroring yet –
should
I?

Replication is the cheap way to do it. Log shipping is the
traditional
way
to do it (and cheap), but your exposure to data loss is larger –
15
minutes +
and it doesn't work well with reindexing. If you want to do
this right
you
need EMC's SRDF or a similar solution like that which costs
millions. The
cost of downtime can make a choice like SRDF economical.

Alastair Jones.