

Re: Best Practices on Writes

Source:

<http://www.tech-archive.net/Archive/SOL-Server/microsoft.public.sqlserver.setup/2005-08/msg00308.html>

- *From:* "Hank Arnold" <rasilon@xxxxxxx>
 - *Date:* Sat, 20 Aug 2005 03:51:22 -0400
-

You're right... It's only a "cluster in the sense that two servers are "connected via software and hardware....."

The software we are using is a highly rated (and not inexpensive) package called LifeKeeper by SteelEye. The mirroring (called Life Keeper Data Replication) is accomplished via a crossover cable on a 10.10.100.x private network. As I understand it, writes are done simultaneously to both drives. The target drive is locked and is not accessible on the target machine.

LifeKeeper itself monitors a "heartbeat" and if it detects that the primary machine is off-line, it will automatically pause the mirror, reverse it and start all the services as well as switch the IP address. This is a protected IP address (each machine has it's own "real" one) that is the one used by all applications to access the server. It accomplishes this in less than 5 minutes.

I have asked them for a "best practices" statement about how they recommend setting it up.

By the way, this is *far* from a "low cost DR" package. It has been working extremely well for over a year and has saved our bacon several times. However, what has changed is that the current servers use standard SCSI (non-RAID) drives and the new servers have RAID arrays with cached adapters.

--
Regards,
Hank Arnold

"Andrew J. Kelly" <sqlmvpnoospam@xxxxxxxxxxxxx> wrote in message news:uUOzSmNpFHA.1416@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx
> Be careful of the term Cluster if it is not an actual Windows Cluster<g>.
> Not sure what software you are using to do this mirroring but I would have
> my doubts as to how effective that is if it is asynchronous. And yes this
> would affect the use of Write caching if the software is getting the data
> directly from disk when it does it's updating to the standby server. If
> the server fails over to the standby you don't have any ability to recover
> those writes in cache. If the process is asynchronous and is at the block
> level I have my doubts it is a viable solution at all. But I have no idea

Re: Best Practices on Writes

> what you have and might be missing something. If you turn off write back
> caching altogether you minimize this risk but you can loose a great deal
> of performance as well depending on your load and the capability of the
> disks. This sounds like a low cost DR solution that may end up costing you
> more if it doesn't work<g>.
>
>
> --
> Andrew J. Kelly SQL MVP
>
>
> "Hank Arnold" <rasilon@xxxxxxx> wrote in message
> <news:exZ6xNNpFHA.4088@xxxxxxxxxxxxxxxxxxxxxxxxxxxx>
>> No. Machine is not clustered. It *will* be part of a High Availability
>> cluster where the backup machine has all the SQL services in manual mode.
>> The primary machine is actively mirroring (asynchronously) its D: and E:
>> drives to the backup machine's equivalent drives. If the primary machine
>> goes off line, the HA software detects this and puts the services into
>> service and switches the mirror roles.
>>
>> In that case, do you think it would make a difference?
>>
>> TIA,
>> Hank Arnold
>>
>> Andrew J. Kelly wrote:
>>> You might want to have a look at this and the associated links within:
>>>
>>> <http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlIObasics.msp>
>>>
>>> Is this machine clustered? I assume not. If not and the cache is
>>> battery backed it is usually safe to enable write back caching. Most
>>> modern controllers will hold the writes until the drive is available
>>> again and write them once it comes back up. The log writes are flushed
>>> thru the cache when written and are the key to ensuring you can recover
>>> in the case of a disk failure. Even if all the data changes were not
>>> written from the cache the database should be able to redo or undo the
>>> changes when it comes back up and goes through recovery. If the disks
>>> are really toast and the whole array gets fried then you have to restore
>>> from last good backups anyway and the cache doesn't even come into play.
>>> Backups are always your best defense for hardware failures. So make
>>> sure you never put your backups on the same array as the data or logs.
>>>
>
>

- **References:**

- ◆ **[Best Practices on Writes](#)**
 - ◇ From: Hank Arnold
- ◆ **[Re: Best Practices on Writes](#)**
 - ◇ From: Andrew J. Kelly
- ◆ **[Re: Best Practices on Writes](#)**
 - ◇ From: Hank Arnold
- ◆ **[Re: Best Practices on Writes](#)**
 - ◇ From: Andrew J. Kelly

- Prev by Date: **[Re: Best Practices on Writes](#)**
- Next by Date: **[/INTEGRATE for SQL 2000 SP4](#)**
- Previous by thread: **[Re: Best Practices on Writes](#)**
- Next by thread: **[Problem about install SQL Server 2000](#)**
- Index(es):
 - ◆ **[Date](#)**
 - ◆ **[Thread](#)**