

Re: Suggestion

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.datawarehouse/2005-07/msg00011.html>

- *From:* "JT" <someone@xxxxxxxxxxxxxx>
 - *Date:* Tue, 5 Jul 2005 13:19:19 -0400
-

Keeping the historical vs. current data partitioned by database is probably a good idea, because it reduces the seek time for the most frequently used data, allows more flexibility in placement of data files on disks, and reduces the size of full database backups, if records for historical years are static. However, if historical 2004 – 1999 data is rarely used, then would there be a cost / maintenance justification to keep it hosted on separate servers? What percentage of the time would Server1 be sitting idle waiting for a user to request a report that includes historical data?

As for spanning historical and current records into one query, read up on "partitioned views" in SQL Server Books Online. For example, the following implements a view of the Customers table, which is split across 3 servers.

```
--Partitioned view as defined on Server1
CREATE VIEW Customers
AS
--Select from local member table
SELECT *
FROM CompanyData.dbo.Customers_33
UNION ALL
--Select from member table on Server2
SELECT *
FROM Server2.CompanyData.dbo.Customers_66
UNION ALL
--Select from member table on Server3
SELECT *
FROM Server3.CompanyData.dbo.Customers_99
```

"Ajay" <Ajay@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message
<news:026E7607-C818-4855-AC9C-9FEF90C623BF@xxxxxxxxxxxxxxxxxxxx>
> Hi,
>
> We have an application that access multiple databases for generating
reports
> and running various queries. These databases are segregated as follows:
>

Re: Suggestion

- > Server1 DB1 1999–2001
- > Server2 DB2 2002–2004
- > Server3 DB3 2005–present
- >
- > Data is sync–ed up using replication. Database grows at an average 4 million
- > records per year. The database is not a Warehouse. Database schema is
- > normalized. Need to keep historical data for generating baseline reports.
- >
- > This approach has various drawbacks including the major one that if any
- > report has to span for the entire duration than data has to retrieved from
- > all databases and then grouped up as summary. This takes lot of time.
- Also,
- > over a period of time, we can't just keep adding more hardware and
- escalate
- > the current problems.
- >
- > So, I am looking at a new database architecture which solves the current
- > problem and is scalable also.
- >
- > Here are the considerations:
- > 1. New architecture should support fast report generation
- > 2. Be scalable to support new report
- > 3. Easy to Manage and Maintain
- >
- > Looking forward to suggestions:
- >
- > thanks,

- **Follow-Ups:**

- ◆ **Re: Suggestion**
◇ From: Ajay

- **References:**

- ◆ **Suggestion**
◇ From: Ajay

- Prev by Date: **Suggestion**
- Next by Date: **Re: Urgen: Visio vs SQL Server**
- Previous by thread: **Suggestion**
- Next by thread: **Re: Suggestion**
- Index(es):
 - ◆ **Date**
 - ◆ **Thread**