

Best Practices Analyzer, multi-core processors, and more

Edited by Nancy Michell

Q I have some questions about the Best Practices Analyzer (BPA) rule concerning the use of Schema Qualified Tables/Views. As per the BPA documentation, this rule doesn't check schema qualification on temporary tables. The BPA report flags references to temp tables created within stored procedures. Should they be qualified? If so, with what schema? I would think the temp tables should be qualified to their owners just like other tables.

A The BPA recommendation to schema-qualify table and view references doesn't really apply on SQL Server 2005, as user-schema separation has fixed the problem that required this practice on earlier versions of SQL Server. Schema-qualification was required to enable query plan reuse by

different users who, in SQL Server 2000, had different default schemas. They could use dbo objects without qualification, but SQL Server had to check the default schemas for the objects first, preventing query plan reuse. Since user-schema separation allows different users to share a default schema and, in the typical case, all database users to access objects in their default schemas, ad-hoc plans with unqualified table and view names will normally be shared and reused across users. Unqualified table and view names inside stored procedures never suffered from this problem in the first place, further weakening the rationale for this practice. BPA in SQL Server 2005 will not include rules like this because it does not have a T-SQL parser. Instead, it will check for improper configuration and security settings.

Figure 1 Find blocking and locks

```
-- script to show blocking and locks
SELECT
  t1.request_session_id AS spid,
  t1.resource_type AS type,
  t1.resource_database_id AS dbid,
  (case resource_type
   WHEN 'OBJECT' THEN object_name(t1.resource_associated_entity_id)
   WHEN 'DATABASE' THEN ''
   ELSE (SELECT object_name(object_id)
        FROM sys.partitions
        WHERE hobt_id=resource_associated_entity_id)
  END) AS objname,
  t1.resource_description AS description,
  t1.request_mode AS mode,
  t1.request_status AS status,
  t2.blocking_session_id
FROM sys.dm_tran_locks AS t1 LEFT OUTER JOIN sys.dm_os_waiting_tasks AS t2
ON t1.lock_owner_address = t2.resource_address
GO
```

Q I'm encountering strange behaviour on SQL Server 2005 SP1. If I delete records on a table and perform inserts on the same table at the same time, the inserts are locked until the delete transaction is committed. Is there a way I can find out what's blocking?

A Yes. You can run the Dynamic Management Views (DMV) script in **Figure 1** to find locks and blockers.

Q In addition to hyperthreading and dual-core technology, chip vendors

are beginning to release processors with additional cores (four, eight, and more). I am considering purchasing a new server with multi-core processors to support a SQL Server 2005 Standard Edition deployment and am curious if, when using a four-core processor, I will only be able to make use of a single physical CPU (since Standard Edition is limited to four CPUs)?

A For the purposes of both licensing and CPU edition support, SQL Server considers only the number of physical sockets/CPUs, regardless of the number of cores on the processor. So, for example, the fact that SQL Server 2005 Standard Edition supports up to four CPUs means it will support four physical CPU sockets, regardless of the number of cores in each one (if you have four physical CPUs with four cores each, your Standard Edition deployment would have 16 logical CPUs to use). Moreover, even though you have 16 cores/logical CPUs, the licensing requires only that you pay for the four physical CPUs, not all 16 cores. For more information on SQL Server and multi-core, see www.microsoft.com/sql/howtobuy/multicore.mspx.

Q I'm about to purchase some new database servers. Should my organisation go with 64-bit or stick to the tried and tested 32-bit versions?

A It depends. This is a very common question now as vendors transform their hardware from x86 to x64. The first question you should ask is what the memory pressures on your database server will be. If your SQL Server instance will consume less than 3Gb of RAM over the life of the hardware (typically three years), then 32-bit x86 is reasonable, provided that will be the direction for other servers (domain controllers, DNS servers, application servers, Web servers, mail servers). If the SQL Server instance will consume more than 16Gb or there will be mul-

In 18 months time it will be almost impossible to source purely x86 hardware

iple instances per server (or cluster), you should seriously consider a move to 64-bit technologies. If there will be more than eight processors, IA64 is the standard recommendation. However, as x64 four- and eight-core CPUs start hitting the streets in quantity, there may be a significant cost argument for going with x64 rather than IA64.

When moving to x64 (or IA64), deploying the same architecture in development, test and performance environments may not have been included in initial cost estimates. But you should also consider whether you want to get stuck with old technology halfway through your hardware life-cycle. In 18 months it will be almost impossible to source purely x86 hardware. If you go with that kind of solution now, upgrading over the next year or two will be expensive as it will be a rip-and-replace scenario. Going for 64-bit hardware now will give you more options over the next few years.

So a good time to make the jump to 64-bit is when you're purchasing new hardware, at the commencement of a core application release cycle, or when you're upgrading from SQL Server 2000 to SQL Server 2005.

Q I have log shipping enabled between two servers. Over the weekend, a hardware problem on the primary server caused log shipping to stop. When I returned to monitor log shipping, I could see that the LS_backup_dbname was running fine despite several hours of downtime.

On the secondary server, however, the copy job appears to be work-

ing fine but the restore job seems to be skipping files and eventually failing over and over again. It runs, but skips every file and then fails. So my time since last restore is now 849 minutes. What I've done in the past in similar situations is simply re-initiate the shipping configuration – that is, I turn it off and turn it on again so I wind up taking a fresh backup and restoring it to the secondary server, and log shipping begins again. Is there a better way to manage this scenario?

A What you are probably seeing is the effect of having missing backup files. As you know, the log shipping backup, copy and restore jobs run independently. Old backup files in the copy location are also removed based on the schedule you specify. If you had a period of time where the restore job was not running, it might have exceeded the interval, which would result in missing backup files.

The log shipping restore job has logic that tries to recover from different types of failures caused by different operational conditions. Basically, it scans the files when errors are encountered and tries to find the correct backup files. If a correct backup file is not found, I suspect something has removed the needed files or possibly you reconfigured the backup part of log shipping?

One solution would be to manually restore up through the transaction log backup files that were inadvertently created by the other backup job. The problem with just copying them to the log shipping copy location is the name conventions recognised by log shipping. But once you have manually restored through that point, the normal restore should pick up and continue.

Q I need to know how to use T-SQL to find the size of total and free physical memory. Is there an easy way to get this information?

A Just run the query in **Figure 2** and you'll get the data you're looking for.

Q I would like to store the sa (system administrator) password in a file encrypted with a symmetric key to make database creation from an application easier. Can I do this?

A First of all, you should really never use sa from an application. Plus, you should use Windows Authentication instead of SQL Server authentication.

If you need to perform highly privileged activities, you should define the minimal permissions necessary to perform the action. Define a database role that contains those permissions. Assign a user to that role. Create a stored procedure with the EXECUTE AS clause that will perform those functions. For an explanation of the Execute As clause, visit www.microsoft.com/uk/technetmagazine/executeas.

Of course, disregarding whether the account in question is the sa account, and bearing in mind that Windows Authentication is not always an option, the more basic question is: how do you securely store a password for a SQL Server login account in such a way that the application that needs to use that password can access it, but unauthorised users or applications can't?

Encrypting the password is a step in the right direction, but it just shifts the problem from protecting the password to protecting the encryption key.

Assuming that this is a Microsoft .NET Framework-based application, look at the Enterprise Library at msdn2.microsoft.com/aa480453.aspx and the recommendations in the patterns & practices section of MSDN at msdn.microsoft.com/practices. The Enterprise Library contains several useful application blocks, including a configuration block, a cryptography block and a data-access block; these can help you to securely store your passwords, but even more, they can help you to manage your data access overall.

If the Enterprise Library isn't an option (maybe you're not using the .NET Framework), familiarise yourself with the CryptoAPI: www.microsoft.com/uk/technetmagazine/cryptography. Here you will find core cryptographic functionality distributed with the OS that can allow you to tie cryptographic keys to a user principal or to a machine, which can simplify key management (which becomes the issue once you've encrypted your password).

Q I want to move my SQL Server 2000 two-node Cluster Server (Active/Active – on 64-bit hardware) from Domain A to Domain B within the same forest. The SQL Server database is in the SAN. The Active Directory® architecture calls for an empty root domain with two child domains (Domain A and Domain B). Domain B is going to be the domain in which all the infrastructure components will be placed. It is a Windows Server 2003 Enterprise environment, with additional domain controllers running Windows 2000.

Is it possible to change the domain membership of the SQL Server Cluster server from Domain A to Domain B?

A For starters, take a look at 'How to Change Domains for a SQL Server 2000 Failover Cluster' at support.microsoft.com/kb/319016 and at 'How to Move a Windows Cluster Server from One Domain to Another' at support.microsoft.com/kb/269196. The same process applies for a multi-instance cluster. The only difference is that you'll have to repeat steps for each instance as appropriate. It's never 'recommended' that you change a cluster's domain, given the complexities involved, but it has been done many times. The only other way would be to do as you stated, rebuild a new instance side-by-side and migrate the data.

Each project is different and a wide variety of tools are used to perform the actual domain-level change (including security considerations, account migrations, sid mappings, and so forth). Often that's the toughest part of the process, not the actual change itself. ■

Figure 2 Get memory

```
With VASummary(Size,Reserved,Free) AS
(SELECT
  Size = VaDump.Size,
  Reserved = SUM(CASE(CONVERT(INT, VaDump.Base)^0)
    WHEN 0 THEN 0 ELSE 1 END),
  Free = SUM(CASE(CONVERT(INT, VaDump.Base)^0)
    WHEN 0 THEN 1 ELSE 0 END)
FROM
(
  SELECT CONVERT(VARBINARY, SUM(region_size_in_bytes))
    AS Size,
    region_allocation_base_address AS Base
  FROM sys.dm_os_virtual_address_dump
  WHERE region_allocation_base_address <> 0x0
  GROUP BY region_allocation_base_address
  UNION
  SELECT CONVERT(VARBINARY, region_size_in_bytes),
    region_allocation_base_address
  FROM sys.dm_os_virtual_address_dump
  WHERE region_allocation_base_address = 0x0
)
AS VaDump
GROUP BY Size)

SELECT SUM(Size*Free)/1024 AS [Total avail mem, KB] ,CAST(MAX(Size) AS INT)/1024
  AS [Max free size, KB]
FROM VASummary
WHERE Free <> 0
```

THANKS TO THE FOLLOWING MICROSOFT IT PROS FOR THEIR TECHNICAL EXPERTISE: Sunil Agarwal, Laurent Banon, Steve Bloom, Chad Boyd, Matt Burr, Shaun Cox, Cindy Gross, Bobby Gulati, Matt Hollingsworth, Arnost Kobylka, Mikhail Shir, Fernando Pessoa Sousa, Stephen Strong and Ramu Veeraraghavan.