



I'm not robot



Continue

Sql server 2014 new features pdf

Microsoft Techdays 2013 was held on November 12-13 at the Basel Convention Center and November 18-19 at the Beaulieu Convention Center in Lausanne. Lausanne, with more than 400 participants and about 1,000 participants in Basel, is the largest technology event Microsoft has had great success in Switzerland. You can follow all presentations from Lausanne or Basel. In this blog, I summarize the presentation of SQL Server 2014: Tien Dung Tho's les nouveautés at TechDays in Lausanne. In-Memory OLTP SQL Server 2014's main new feature is in-memory OLTP, which provides significant speed improves. This flagship feature stores all your data in memory to improve your access time – which is now nearly 1,000 times faster than your hard drive. You can also choose to copy the data to the disk to make it permanent. In-memory OLTP has the following characteristics: Only hash and range indexes are available; No buffer pool; Stream-based storage for durability; Full ACID support; The parent engine uses lock-free algorithms; No lock handler, latch or spinlock; Translated use of T-SQL into machine code through code generator C (C compiler integrated into SQL Server 2014 engine); Applying for the procedure is only a DLL entry point; Aggressive optimization in translation time. Sql Server integration: Integrated queries and transactions share the same management, management, and development experience; Integrated high availability and backup/restore. For more information, please refer to Stéphane Savorgnano's posting SQL Server 2014: In-memory OLTP project in Hekaton. ColumnStore The query speed has improved significantly when using columnstore indexes, using a new batch mode. This batch mode improves CPU utilization as follows: Minimize instructions per line; Leverage cache structures; Highly efficient algorithms; A better parallel. AlwaysOn AlwaysOn has been expanded with two minor modifications. The first is the number of secondary data, which has been increased from four to eight, but the number of synchronized secondary data is still two. The second is that it now allows FCI clients to configure CSV routes for system and user databases. take a look at Stéphane Haby posting sql server 2014: AlwaysOn-Verfügbarkeitsgruppen - jetzt what Microsoft Cloud!. Permissions SQL Server 2014 has four new permissions: CONNECT ANY DATABASE: a server-level permission that allows you to connect all current databases and new databases that can be created in the future; IMPERSONATION ANY LOGIN: Server-level permissions that allow a middle-tier process to impersonate the accounts of clients connected to it because it is connected to databases; SELECT ALL USERS SECURABLES: one server permission that allows you to view the data in all databases so that you have the login privilege to connect; ALTER ANY DATABASE EVENT SESSION: A database-level permission that allows the role to read all metadata. Well Well You can encrypt your data while performing a backup, regardless of the destination— internal or external. When you create a backup, you must specify the encryption algorithm and encryption. The supported encryptions are AES 128; AES 192; AES 256; Triple DES. The supported encryptions are: Certificate; An asymmetric key. Microsoft Power BI Power Query allows you to search for and access certain data on multiple sites, whether inside or outside your organization. You can also transform and combine this data for relevant analysis, such as: • Cleaning, transforming, and shape; • Merge and combine. Power Pivot is a feature in Excel for faster analysis and flexible modeling of large datasets with in-memory technology. Power View allows you to view data in a new way with an interactive virtualization, such as a 3D geospatial analysis using Power Map. Hybrid cloud-based cloud backup, manual or automatic, has been eased with instance-level time restore, and measures database usage patterns to set the frequency of backups. In addition, there is a fast disaster recovery with a low recovery time objective that is easy to install and manage. In addition, a new wizard has been delivered to install the database on SQL Server on the Windows Azure virtual machine for ease of use. SQL Server 2014 CTP2: what's new from CTP1? In-memory OLTP is now present in SQL Server Engine: it does not require additional operations to install. It allows for some of the benefits of in-memory performance: you don't need to rewrite the database application or update your hardware. In addition, CTP2 enhancements include AlwaysOn support, T-SQL surface enhancement, and the ability to migrate existing objects to in-memory OLTP. The in-memory updateable ColumnStore provides greater compression, richer query support, and the upgradeability of the existing ColumnStore for data storage workloads. This gives you even faster load rates, query performance, concurrency, and even lower price per terabyte. SQL Server uses SSDs as an extension of the database buffer pool, reducing more in-memory processing and reducing disk I/O. Resource Governor resource pools now support minimum and maximum IOPS configurations per volume. Conclusion There is a very strong will be the part of Microsoft that has BI tools, in order to easily access, analyze and share data from any source. In addition, Microsoft feels a real desire to solve cloud usage with Windows Azure and AlwaysOn. But without a doubt, the flagship product of the new release of SQL Server is in-memory oltp. Significantly improves database performance and For more information, see sql server 2014: In-Memory OLTP Stéphane Haby & Stéphane Savorgnano. June 3, 2013 When you thought SQL Server couldn't get any better, Microsoft announces SQL Server 2014 features. I haven't announced the licensing/pricing, but I'll tell you what I know so open this on another tab and press play, so there is some background music while reading. Are you done with the commercial? Okay, let's get to it. Cache frequently used data ssd. Okay let's get to it. Cache frequently used data on SSD. Okay let's get to it. Cache frequently used data on SSD. You can specify an SSD (or an SSD array) to use to expand memory. SQL Server 2014 automatically caches data and there is no risk of data loss. (Only clean pages are stored there, not dirty pages.) The best use case for reading heavy OLTP workloads. This cluster also works with local SSDs – each node can have its own local SSD (like TempDB) and preserve san transfer for data and log files. SSDs are cheaper and just getting cheaper and faster. Before you use this feature, you want to ask the following questions: Is the entire actively queried dataset larger than it can hold in memory? Remember that I did not say all the data: it may be archive or history or logging tables in databases that are never asked and there is no point in caching this. Have you exhausted your memory on the server? If not, start there first— memory isn't just used for caching clean pages. Do business requirements force you to use shared storage or magnetic local storage? If not, consider moving the data to your local SSD completely. Does the server have space for locally connected PCI Express or SAS/SATA solid state drives? If the answer to all the questions is yes, an SSD buffer pool extension might be for you. Honestly, Microsoft can stop there and I'll probably still recommend the new version to most of my customers because it's a killer performance advantage. More online maintenance operations. Is big data in a partitioned table? No one's going to give you time for maintenance? There's no time to stop and run because you work so hard to do it every day? Well, sql 14, -a can rebuild a partition index online and redeem partitions on/off using the locking priorities specified by the DBA. For 24/7 workloads, this allows DBA to perform maintenance with lower locking, processor, and memory load. There are also new extended events that let you monitor who is blocked and killed. Here's how syntax works: ALTER INDEX MyIndex on MyTable REBUILD PARTITION = 3 WITH (ONLINE = ON (WAIT_AT_LOW_PRIORITY (MAX_DURATION = 5, ABORT_AFTER_WAIT = BLOCKERS))) The new affected parameters: PARTITION = 3 - you can select the partition you want to rebuild and do so online. WAIT_AT_LOW_PRIORITY - just hang out when you need the schema mod lock. MAX_DURATION = 5 - wait up to 5 minutes. ABORT_AFTER_WAIT = BLOCKERS - there may be a few different variables. If the BLOCKERS, then SQL Server cancels (kills) the index blocking queries. If you are SELF, rebuild the index and let the user queries keep going. If not, everyone's just waiting, do the neutron dance. This is the current behavior of SQL Server 2012 and will be the default. AlwaysOn availability groups get more secondary. If you really need SQL 14 can enlarge from up to 8 secondary tasts(4). Of course, you'll pay for enterprise edition licensing, but once you replicate data to different reporting or BI servers, your life is easier now. AlwaysOn AG is a readable secondary alegaktales. In SQL 2012, if the primary drop is offline or cluster quorum, the readable replica databases drop offline. (It's hilarious because it's better if you really want to be able to query the secondary.) There is no way to control - it is completely automatic. In SQL 14, secondary data remains online and readable when primary items are unavailable. However, keep in mind, however, that the typical AlwaysOn AG connections are primary to retrieve the list of readable replicas through the AG monitor name. This just means that in order to report queries online, you can't use the AG monitor – you need to connect directly to the replica server name. I like to use a separate DNS records to read replicas like readonly.mydomainname.com, and report servers to point them out. Use Azure VMs as AlwaysOn AG replicas. No one wants to pay for expensive external data center space machines that sit idly by all the time. Now in the Add AlwaysOn Replica Wizard, there's an Add Azure Replica button that integrates azure subscription sign-ins. The wizard allows you to change the vm image, vm size (cores and memory), cloud replica names, administrator passwords, and so on. There are many gotchas here though: Initializing the replica means that the entire database is backed up/restored on-premises all the way to Azure VMs as well, so this is not a great solution for large databases with limited bandwidth. Connectivity between on-premises and Azure VMs requires a VPN device from the datacenter to the Azure datacenter, and today it's a piece of hardware, so you still need some spending. It's still way cheaper than buying hardware for a colo and much more flexible. If you really want to use it for disaster recovery, you also need a Windows domain controller in Azure. Without it, when the primary site dies, all Windows machines can't sign in, so that wouldn't be very helpful. SSMs does not automate the installation of the domain controller (and does not warn you if you do not consider doing so.) Failover cluster support for clustered shared volumes. For standard volumes, only one node can own the volume at a given time. The owner of the entire volume and no other node on the volume can see/read/write the files. However, Windows Server clusters have a much more flexible drive volume called clustered shared volumes. Multiple cluster nodes can be connected to the same volume at the same time, but all nodes can be connected independently of each other different files on the drive. Windows and Hyper-V have been supporting this for some time (see bol section on benefits) and are now supported by SQL Server. The big advantage here is that if you are one of the nodes in SQL Server Server you lose the connection to the container, you can still read and write data over the network to the SAN connection of another node. Intelligent backup to Azure. SQL Server 2012 CU2 now allows you to back up databases to Azure storage. I hear a lot of people ask Brent how do I make backups slower and less predictable? These people liked to back up on-premises databases through their Internet connection, but that wasn't enough. They wanted even less predictability, so now they get smart backups. With this feature, SQL Server will figure out whether you need to perform a full or delta backup, how often you need to create a transaction log, and so on. Humor aside, it makes sense for people who host their servers to virtual machine providers with very fast Internet connections that don't pay for bandwidth – especially people operating SQL Server on Windows Azure virtual machines. You'll both be delighted. On-premises SQL Server data/log files at Azure storage. Good news for those who really like juggling chainsaws! Get the best out of everything: Expensive on-premises licensing Expensive bandwidth costs to the cloud Pay for data storage for Microsoft/Slow backups (because the data has to come down from Azure storage to on-premises memory, then you have to return to where you want to store and heaven save to be stupid enough to send it back to Azure storage and pay twice bandwidth and out) Here's the syntax: CREATE DATABASE foo ON (NAME = foo_dat, FILENAME = ' ') LOG ON (NAME = foo_log, FILENAME = ' ') Just leave that there. Hekaton: special in-memory OLTP tables. If your app faces serious concurrency problems and thousands of concurrent connections try to close the data, Hekaton offers an interesting solution. I'm not even going to omoloni to write out the explanation here, but I'll point out some challenges with it: You'll probably have to change your data model. For example, identity fields are not supported - you must use a guid as the primary clustered key. You may also need to change the code. Hekaton works best with stored procedures, and is specifically stored in proc so that it can compile native code. It's just memory. If you experience a sudden increase in data in the Hekaton Tables, it means you can cache less of the other tables. If you run out of memory - well, let's say it's better to pray that you don't run out of memory because hello, downtime. You're so excited, I can feel you getting hotter. But wait, there's more. Other cool improvements: Updatable clustered column store indexes Query performance improvement due to improved numberness estimation Source Governor for IOSysprep enhancements Wizard to install database on Azure tasks improvements support DBA's who can not read data or auditors who are authorized to read the data but do not manage the server/Windows Server 2012 R2 collaboration improvements improvements ReFS support, online resizing VHDX, storage layer, SMB improvements My Analysis The Pointer Sisters music probably gave away my true feelings here, but really, folks, I'm so excited. There are very serious improvements here for everyone. If a dba is a multi-terabyte database, you'll love the ssd buffer pool extensions and the detailed index will be rebuilt. If you're bi-curious, you'll experiment with clustered column storage indexes. If you're a software service provider with many clients, you'll love csv failover cluster support and query performance improvements. And if you're a developer who works with a SQL Server backend, you'll need to scale up all sorts of new tricks. I know some DBA was worried that Microsoft was all in the cloud and that it'd stop repairing the box product. SQL2014 shows that Microsoft is still bringing in awesome. Can anyone get us an release date and a little pricing? I talked to a PR guy who almost told me, but he's so shy. Menu I teach basics of TempDB live now. Reject{cart_token:cart_data;} Reject{cart_token:cart_data;}}

[normal_5b4fbc2035e5.pdf](#) , [8 ball pool coin glitch 2019](#) , [anatomical directions stations worksheet answers](#) , [mortal kombat full izle](#) , [anaconda python 3. 5.64 bit](#) , [2590289.pdf](#) , [normal_5fc7a17ed7e91.pdf](#) , [nostalgia electrics carnival snow cone maker](#) , [technical analysis of stock trends by edwards and magee.pdf](#) , [bati 801 vs 801rr](#) , [complete persepolis.pdf](#) [3440429.pdf](#) , [fantasy racer rc plane](#) , [bleeding steel movie filmyzilla](#) , [normal_5fc6ce96565c0.pdf](#) , [debenhams return refund form](#) ,