



SIOS WHITE PAPER

Calculating Costs for Microsoft SQL Server

On-Premises vs. Cloud




In partnership with



Denny Cherry
& Associates Consulting

us.sios.com



One of the items that invokes the most nightmares when moving to the cloud, is the cost of big-ticket items such as Microsoft SQL Server. When done incorrectly the costs of putting a database server in the cloud can be daunting, but costs can be controlled with some understanding of the Microsoft SQL Server licensing process, and how you can use some of the requirements to your advantage.

Purchasing Licenses vs. Renting Licenses

One of the first questions that you should be asking yourself before moving database servers to the cloud, is: should I rent the Microsoft SQL Server license from the cloud provider or should I purchase the Microsoft SQL Server license and bring my own license to my Virtual Machine. All three major cloud providers (Microsoft Azure, Amazon Web Services, and Google Cloud Platform) will happily rent you your Microsoft SQL Server license. You do not own the license. However, if you already do own the needed number of Microsoft SQL Server core licenses (and Software Assurance for those cores) the discussion quickly switches to whether your company should continue to pay for Software Assurance for the licenses, or begin renting the licenses instead. Although some organizations fear moving to the cloud means losing some control over their infrastructure, the cloud enables new opportunities for the technology team to focus on core business initiatives with cost savings.

Software Assurance

When using Software Assurance licenses on a virtual Microsoft SQL Server it is important to remember that in order to use that license in the cloud the license must be kept current. In the same way, it needs to be kept current in order to continue to virtualize Microsoft SQL Server in an on-premises environment. The reason for this comes down to the license mobility right that Software Assurance gives you. This right gives you the ability to move the Microsoft SQL Server license from one set of CPUs to another (running on a different physical server) whenever you need to. In a physical virtual environment, this covers your requirements for VMware's vMotion or Microsoft Hyper-V's Live Migration. In the cloud, this also covers your need to move the license between physical machines in the event the host computer fails or restarted.

In a clustered environment, Software Assurance also gives you an important license right-- the ability to run a secondary node without paying for the license for that secondary node. This means that when using Microsoft SQL Server in a failover cluster environment, only the currently active node needs to be licensed, assuming that you own the license and Software Assurance for that node. In the event of that node's failure (or any failover event) the license will be automatically transferred to the node, which was passive, and is now the active node. Without Software Assurance, this right does not exist, and both nodes would need to be licensed.

Renting Licenses

Renting the Microsoft SQL Server license comes with some advantages as well. The biggest advantage is that there is no cash outlay required when renting Microsoft SQL Server licenses, as the licenses are rented either through the cloud provider with the cost of the Virtual Machine or rented through a third-party vendor.

The second big advantage of renting the SQL Server license is the ability to scale the VMs up or down, which will adjust the Microsoft SQL Server license cost accordingly. If the VM that is hosting the Microsoft SQL Server service is scaled down due to a reduction in load, the Microsoft SQL Server license costs are reduced as well. With software assurance costs, you cannot reduce the Software Assurance costs unless you are sure that you won't be scaling the virtual machine back up in the future. The reason for this is that once you stop paying to renew the Software Assurance benefits for a set of cores, you don't have the ability to restart the Software Assurance license. So, if the needs of the application were to grow again, you'd need to buy additional Microsoft SQL Server licenses as well as the Software Assurance license for them while when renting the license. You can then simply increase or decrease the number of cores that you need. Assuming the license rental fees are calculated on an hourly basis and the end of the hours of the scale down the license fees would be reduced.

SQL Server Edition Selection

One of the places where large cost savings can be achieved is by using SQL Server Standard Edition over SQL Server Enterprise Edition. By moving the database server edition to SQL Server Standard Edition and moving the database high availability solution to Windows Clustering we can maintain the high availability of the database that is needed, while keeping the instance costs under control. Using SQL Server Clustering (and Windows Clustering in general) requires that a shared storage configuration be enabled. While no cloud vendor has a native storage clustering solution, we can achieve the needed configuration by using SIOS DataKeeper Enterprise Edition. SIOS DataKeeper Enterprise Edition will allow us to cluster the storage that is presented to the virtual machines which are members of the SQL Server cluster, and it will present that storage to the Windows operating system in such a way that Windows sees the storage as shared storage. Replication of the data between the active and passive nodes can be configured for synchronous data movement or asynchronous data movement depending on the network topology design. In either solution SQL Server Standard Edition can be used which will greatly reduce the SQL Server licensing costs for the solution.

For solutions that require SQL Server Enterprise Edition combined with Failover Clustering, SIOS DataKeeper Enterprise Edition will be the storage replication mechanism of choice over the Native Windows S2D solution.

Virtual Machine Sizing

Sizing of workloads when moving to the cloud is one place where companies can save large amounts of money by right sizing workloads.

Production Nodes

The node of the production environment - which is the active node - is the most expensive node. This is because the Microsoft SQL Server service is active on this node, and it currently services queries. When sizing this node for a move from an on-premises environment to a cloud environment there are going to be a couple of things to take into account.

1. Adding CPU resources in the cloud takes seconds¹
2. Cloud CPUs may be faster than on-prem CPUs²

When looking at CPU sizing of a virtual machine it is important to remember that there should be a thought shift from the days of the on-premises world. In the on-premises world we would need to design the server for what the CPU load of the server needed to be 3-4 years in the future, when the server would be reaching the end of its lifetime. In the cloud world, we do not need to plan our sizing with that longevity in mind. Adding CPU in an on-premises environment takes coordination to scale the virtualization farm up, and possibly buy additional resources as the on-premises servers could only be scaled up so far. Eventually you would hit the top end of the number of CPUs that could be put into the server. If the Microsoft SQL Server is running on physical hardware, scaling the server up can be even harder. The server may be out of CPU sockets, or if more CPU sockets are available that will probably activate additional NUMA nodes, which will require the purchase of additional RAM as well as CPUs in order to use new CPUs. And none of this takes into account CPU price. CPUs aren't cheap to purchase with the newest processors costing tens of thousands of dollars per physical CPU.

In the cloud world, we don't have to worry about any of this. In the cloud world we simply select the machine in the cloud vendor's portal and select the resize option, select the size that you need and submit the change. The virtual machine will restart and the virtual machine will be resized when it comes back online. Total downtime to the machine to resize it, 30-60 seconds.

When it comes to the scale of servers that are available, you'll be hard pressed to find physical VMs that can run at the scale of cloud servers. Depending on the cloud platform selected, cloud-based virtual machines can scale to over 400 cores in a single VM.³

The second item to review is the speed of the CPUs that the cloud vendor is able to provide. On-premises CPUs are probably 1-2 years old at this time and will typically be slower than the CPUs that cloud vendors are able to provide for virtual machines. Additional on-premises virtualization farms are typically over-provisioned, thereby slowing down the database server CPUs; while cloud providers do not over-provision the host servers allowing the guest OS to have full access to the CPU resources that are available.

With these changes in the environment fewer CPUs on the cloud-based server may be needed than on the physical server, especially when taking into account that the virtual machine doesn't need to be sized based on the potential size needed in 3-4 years, but instead it can be sized based on the needed size of the virtual machine and its expected load for tomorrow. This allows the CPUs of the virtual machine to be run at much higher CPU loads as we don't have to account for overhead for months or years from today.


Production Secondary Nodes

Second nodes in the production site can be scaled at a smaller size than the production

¹ Different cloud vendors bill rented licenses at different intervals with some licenses being rented on a per minute basis, and others being rented at a per hour basis. Consult your cloud vendor to see how their rental licenses are billed. Hourly billing is only used here as an example.

² As of April 2021.

³ Each cloud vendor has different top end Virtual Machine sizing options. Check with your cloud vendor to ensure they are a Virtual Machine which will satisfy the sizing requirements of your workload.



servers. These servers can be scaled at a size that will allow the application to run acceptably while this server is running the production workload. Then the service can either be failed back to the larger server, or if that server has failed, the server that is running the workload can be resized as needed to run the production workload. Again, we're taking advantage of the fact that the cloud gives us the ability to scale resources which we simply didn't have in an on-premises world. When sizing down the secondary nodes of a SQL Server environment, no impact should be had on SIOS DataKeeper Enterprise Edition as the secondary node will have very minimal CPU and memory requirements.

Disaster Recovery Nodes

With disaster recovery nodes, we can save additional money. Disaster Recovery nodes can be sized at a much smaller size than the production nodes. Given that most disaster recovery environments are going to use an asynchronous data replication topology and that this requires human interaction to initiate the failover to the DR platform, the database server can be scaled up to the production size in the event of a disaster event when the production services need to be failed over to the disaster recovery site.

Licensing of the disaster recovery site is another place at we can save money on our Microsoft SQL Server licensing. If using licenses that have been purchased a smaller number of licenses are needed for SQL Server for the disaster recovery environment. In the event of a production failure the licenses could be "transferred" to the disaster recovery machine when that machine is scaled up.

If renting licenses for the disaster recovery server, the additional license cost wouldn't come into play until that server is scaled up. And when that happens the production servers are offline so there should be no license rental cost for those Virtual Machines while they are offline.

For disaster recovery nodes within a cloud environment which are going to be scaled up before the servers are needed, the servers can be built and run at an extremely small size. The only real workload that is going to be running on these servers until the failover event is the workload of replicating data to the machine. As long as the machine is sized large enough to handle this data replication workload, there's no need to size it bigger. With Microsoft SQL Server having a minimum license count of 4 cores, there's no real reason to have a disaster recovery server with less than 4 cores as you are paying for 4 cores of Microsoft SQL Server licensing for the machine.

Reserved Pricing

Each major public cloud vendor offers some flavor of reserved pricing . With reserved pricing you guarantee that you will use the cloud resource for a period of time, and due to that, the cloud provider discounts the pricing for that resource. The amount of these discounts vary by cloud platform and they can vary based on the size of the resource that you purchase the reservation for, as well as the amount of time that you guarantee the usage for. These reservations can include a 30% discount on the cost of the cloud resource, a discount which can be realized right away.

Virtual Machine Sizing

When selecting the size for the production server, and reviewing the workload on the existing server, review the peaks of the workload on the server. Understand what is triggering those peaks in workload to happen and understand what will happen if those peaks flatten out into little plateaus as the amount of CPU available of the server is reduced.

One of the traps that IT professionals fall into is looking at the average usage of servers when sizing those servers for a cloud migration. By looking at the averages instead of the peaks you'll have a skewed view of the server's workload. For example, a server is used by a team that works 12 hours a day on weekdays and they push that server to an average workload of 70%.

	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Worktime	0	70	70	70	70	70	0
Non-Worktime	0	0	0	0	0	0	0

Figure 1: Workload Averages

When you look at the average for the entire work week, they are only using 35% of the servers' resources. When you look at the average for the entire 7-day week, then the server is only using ~25% of the servers' resources. If we then built our cloud-based server on the assumption that we only needed a server with 25% of the resource of the on-premises server, the migration would be a failure as the server would not have the resources available.

Storage Pricing


Cloud-based storage pricing is significantly less expensive than on-premises SSD-based storage. While on-premises SAN storage can typically cost several thousand dollars (US) per terabyte, cloud-based storage is much, much cheaper. While each of the major cloud providers price storage differently, Enterprise SSD-based cloud storage is currently costing about \$150 (US) per terabyte, with each of the major cloud vendors having roughly the same price for storage.

Employee Locality

With an on-premises environment the company has to have employees physically in the same locations as their data centers, or at least close, with close being a couple of hours drive away at most. This can be limiting to a lot of companies, especially smaller organizations that want to be able to grow into a new region, but first they need to hire staff to manage the servers that are going to be used by the customer in that region.

By shifting to the cloud, the company no longer needs to physically put employees in the same location as the servers. Virtual servers can be managed from anywhere in the world, so the company can focus on hiring the employees that will best service their needs instead of hiring the employees that live in the right part of the country or the world.

This may very well mean the companies that are used to hiring employees in a certain major metro area can look at hiring employees outside of that metro area, either in a more rural area



or a totally different metro region that has less expensive employees. As long as the person managing the cloud resources has decently faster internet access, they can successfully manage cloud resources for the company.

And because the company does not need to physically manage the servers anymore, this means no late night drives out to the data center to replace hard drives, and this is all managed by the cloud provider, leading to reduced cost for the company, and a better life experience for the employee.

Reducing Cost

Moving resources to the cloud can be an expensive operation. But it does not have to be. There are a variety of ways to reduce cost which are simply overlooked by many companies as they begin their cloud migration process.

When the cloud migration process is done correctly, servers are sized correctly and reservations are put in place to reduce costs further, moving to the cloud can be cheaper than purchasing new servers and doing a full server refresh.

How do I get started with SIOS?

Check out the resources below to get started with SIOS DataKeeper:

- To learn more about SIOS visit - <https://us.sios.com/products/windows/datakeeper-cluster/>
- Schedule time with a SIOS Solution Architect - <https://us.sios.com/demo-request/>
- Request a Trial Copy of SIOS High Availability Clustering Software for Windows or Linux - <https://us.sios.com/san-sanless-clusters/free-trial-evaluation-san-sanless-clusters/>
- Contact Sales: 1.617.245.6955 or sales@us.sios.com



SIOS Technology Corp.
155 Bovet Road, Suite 476
San Mateo, CA 94402
Tel: 650-645-7000

info@us.sios.com
<https://us.sios.com>

© 2023 SIOS Technology Corp. All rights reserved. SIOS, SIOS Technology, SIOS DataKeeper, SIOS LifeKeeper, SIOS Protection Suite and associated logos are registered trademarks or trademarks of SIOS Technology Corp. and/or its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. WP-1035-B