



BRENT OZAR
UNLIMITED®

When is the Cloud Right for You?

Agenda

Defining “the cloud”

Separating responsibilities

Server requirements

Potential bottlenecks

Offloading opportunities

What is “the cloud”?

The cloud is:
Someone else's servers
In a pool of **shared resources**
Where we don't have **total control**.

A cloud is probably virtual

Similar considerations as virtualization.

- Shared resources.
- Measure HA and DR.

Less control over features and functionality.

Examples of clouds

Public cloud

- AWS
- Azure
- Rackspace

Private cloud

- Large VM farm
- Internal IaaS (OpenStack, Eucalyptus)

Agenda

Defining “the cloud”

Separating responsibilities

Server requirements

Potential bottlenecks

Offloading opportunities

Separating Responsibilities

What is the cloud responsible for?

Your cloud provider has limited responsibility.

They provide:

- Facilities
 - Power
 - Cooling
- Infrastructure
 - Physical machines
 - Network equipment
 - Virtualization infrastructure

What if hardware breaks?

That's not your problem.
You're renting it.

What are you responsible for?

Everything
else.

Does an SLA apply
to me?

We don't need an SLA, it's the cloud!



13

© SQLIntersection. All rights reserved.
<http://www.SQLIntersection.com>

If you don't know...

SLA: Service Level Agreement

RPO: how much data can we lose?
(Recovery Point Objective)

RTO: how long do we have to get it back?
(Recovery Time Objective)

This can also be related to performance:

- Users must authenticate in < 15ms
- The shopping cart must display in 250ms
- Any catalog page must render in 400ms
- Reports must perform < 100,000 on the OLTP system.



14

© SQLIntersection. All rights reserved.
<http://www.SQLIntersection.com>

It sounds like we really do need an SLA!



© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

15

If you don't have an SLA, get one!

SLAs will govern everything you do.

- What happens if your provider loses power?
- What happens if the eastern seaboard loses power?
- What happens if someone drops the storage offline?

Losing control of the underlying architecture will make you painfully aware that these things can happen.



© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

16

Helpful SLA tips

Head over to <http://brentozar.com/go/fail> to get our HA/DR worksheet.

Talk to the users to put together a performance SLA.

- Start with 5 queries or pages.
- Time the whole application stack.
- You may need new monitoring tools to support this.

Agenda

Defining “the cloud”

Separating responsibilities

Server requirements

Potential bottlenecks

Offloading opportunities

Server Requirements

Current hardware?

What does the current hardware look like:

- CPU count & speed
- Memory
- Network
- Disks

CPU

Examine CPU usage over time.

Carefully monitor the source of CPU consumption.

- Try to tune for CPU.
- Tune indexes. Tune queries.
- Offload work.

Do you *really* need extra CPU?

- More CPU costs more money **per hour**.

CPU Costs

AWS – \$0.15 / CPU / hour

Azure – between \$0.16 / CPU / hour
and \$0.30 / CPU / hour

Rackspace – between \$0.17 / CPU / hour
and \$0.34 / CPU / hour

Identifying CPU Consumers

Find your top CPU consumers with sp_BlitzCache™

<http://brentozar.com/blitzcache>

Identify missing indexes with sp_BlitzIndex®

<http://brentozar.com/blitzindex>



23

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

Memory

How much data do you have today?

How much of that data are you querying right now?

How fast is your database growing?

What else is going to run on this server?

The correct answer is “nothing”.

Do you *really* need Enterprise Edition?



24

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

What about disks?

I/O is difficult to measure

What's an IOP?

- Amazon says it's 16KB.
Unless it's an SSD. Those use 256KB IOPS.
- Azure is "one IOP is one IOP".
- Rackspace is eerily silent.
- SQL Server does 8KB or 64KB or, sometimes, 2MB IOPS.

Check out our memory sizing post for I/O measuring techniques.

- <http://brentozar.com/go/memory>
- Remember your requirements, they help you pick the size of your cloud server.

Everything has a measurable cost

CPU is by the hour.

Memory increases require CPU increases.

Disk costs along multiple dimensions:

- Capacity.
- Throughput.
- IOPS.

Network has ingress/egress charges.

Plan wisely, monitor carefully

Every provider lets you monitor and alert on usage conditions. Use them.



How should you size cloud servers?

Measured guesswork

1. Evaluate your current utilization.
2. Extrapolate requirements.
3. Decide on your tolerance for overage.
4. Size appropriately.

Two ways to move

Method 1: Forklift

Typically start with a 1:1 match and tune later.

You should still tune first, but...

- Requires additional time and effort.
- No guarantee of fixing any problems.
But you probably will.

Tune after you move to the cloud.

Method 2: Redesign

Set aside an architecture period.

Redesign your application to work with the cloud.

Downsides:

- Time consuming.
- Still no guarantee that it will work.
- You'll have to tune after you get there.

One last thing about servers...

Don't think of it as a server

What do you have to do to get a new physical server?



Don't think of it as a server

What do you have to do to get a new physical server?

What do you have to do to get a new cloud server?



37

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

```
jeremiah@heavy-messing: ~  
jeremiah@heavy-messing ~$ aws ec2 run-instances \  
> --image-id ami-efef19fdf \  
> --count 1 \  
> --instance-type m3.2xlarge \  
> --security-group-ids sg-5a41d76a \  
> --subnet-id subnet-ddcc7474b6 \  
> --ebs-optimized
```

New SQL Server 2014 on Server 2012R2
With 8 cores and 30GB of memory
In my SQL Server network
With RDP access from my office
And preferential access to disks
...
...
Then wait for sysprep to finish

Scripting helps ensure success

Script SQL Server deployment and configuration.

Deploy new environments rapidly.

- Test changes
- Stand up new customers
- Disaster recovery

Sometimes, it's easier to rebuild a system from scratch than it is to troubleshoot.



39

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

Three Deployment Options

PaaS: SQL as a Service	IaaS Standard: SQL in an EC2 VM	IaaS Enterprise: SQL in an EC2 VM
Zero data loss Automatic failover Reduced features, datatypes, control	1 hour data loss 1 hour manual failover Same SQL Server you already know and love but less features	1 minute data loss 1 minute manual failover Data compression Snapshots
\$4.4k per month per server	\$4k per month per server	\$56k up front, \$1.5k per mo

* Per server means "per mirrored pair of servers"



40

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

Agenda

Defining “the cloud”

Separating responsibilities

Server requirements

Potential bottlenecks

Offloading opportunities

Potential Bottlenecks

Measure your top waits!

This will point you to your SQL Server bottlenecks.

<http://www.brentozar.com/go/waits>

Figure out when you have your worst waits:

- Storage waits after hours may not be bad.
- There are always mitigations for problems you find.

Don't get the fastest storage you can – try to mitigate storage issues first.

But what about disks?!?

Storage is terrible.

Storage is usually the slowest part of SQL Server.

The cloud is no exception.

Plan for narrow throughput to storage.

- Storage throughput is often limited by server size.
- Even fast servers typically have 1GbE.
That's 125 MB/s
- If you need faster, you're going to pay for it.

But what about IOPS?!?

If you really need to know IOPS...

Break out perfmon and measure:

- Physical Disk – Disk Reads/sec
- Physical Disk – Disk Writes/sec

This is about as accurate as it's going to get.

Make some guesses from there about your IOPS requirements.

Accurately measuring disks

Look at dm_io_virtual_file_stats

- Average out your disk throughput.
- Look for high stalls.

Record metrics with perfmon

- Physical Disk – Disk Reads/sec
- Physical Disk – Disk Writes/sec
- Physical Disk – Avg. Disk sec/Read
- Physical Disk – Avg. Disk sec/Write

Use cloud provider metrics

Record your key metrics, as the provider sees them:

- Data throughput to disks.
- Disk latency.
- IOPS

Yes, you collected this at the OS level, but also collect this at the infrastructure level.

What if the server won't fit in the cloud?

Hypothetically...

You have a SQL Server that has a resource intensive workload.

- You have more than 32 CPUs.
- You have an all SSD SAN.

Management is all in on the cloud.

What do you do?

- a) Ragequit
- b) Sob quietly until the problem goes away
- c) Deal with it



Dealing with it: offloading work

Do less work in SQL Server.

Move functionality out of SQL Server:

- Cache common queries.
- Move search to external services.
- Sort in the application.
- Perform paging anywhere else.

Split the workload

- Report on server A
- Transactional work on server B
- Single sign on/user profiles on server C



53

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

Your cloud provider can help

Let someone else host other applications.

- Need caching? Spin up an Amazon ElastiCache or Azure Cache server.
- Need bulk data processing? HDInsight or Elastic MapReduce

Focus on your core skills and functionality.

You can always create specialized VMs later.



54

© SQLintersection. All rights reserved.
<http://www.SQLintersection.com>

Common Blockers

What stops a move to the cloud?

Unwillingness to change.

Unrealistic expectations.

Lack of planning.

What should I measure?

Waits: <http://brentozar.com/go/waits>

Memory & I/O: <http://brentozar.com/go/memory>

Worst Queries: <http://brentozar.com/blitzcache>

Wrapping Up

In summation...

There's no hard and fast "This will work, that won't"

Measure your systems:

- Know load patterns.
- Understand SLAs.
- Understand the deployment platform.