



Global Knowledge®

Expert Reference Series of White Papers

Top Four
Considerations
When Choosing
Your SoftLayer Data
Center Location

Top Four Considerations When Choosing Your SoftLayer Data Center Location

John Hales, Global Knowledge VMware, SDN, and SoftLayer Instructor, A+, Network+, CTT+, MCSE, MCDBA, MOUS, VCP, VCAP, VCI, EMCSA

Introduction

In teaching SoftLayer classes around the world, one of the questions frequently asked and discussed is how to choose a SoftLayer data center. Several variables go into that decision, which is the topic of this white paper. These reasons broadly fall into two categories, namely business ones such as cost, and technical ones such as latency or feature availability. While there is no single answer (otherwise there would only ever need to be one data center), there are guiding principles to help make optimum selections for a given set of needs.

Before we begin, it is helpful to know the data centers that SoftLayer has available worldwide. As of the writing of this paper, the available data centers are the following:

- North America:
 - Dallas (4)
 - Houston
 - Montréal
 - Querétaro (Mexico City area)
 - San Jose
 - Seattle
 - Toronto
 - Washington, DC
- Europe:
 - Amsterdam
 - Frankfurt
 - London
 - Paris
- Asia:
 - Hong Kong
 - Singapore
 - Tokyo
- Australia:
 - Melbourne
 - Sydney

Of course new data centers come online frequently. You can always access the most up-to-date list by going to www.softlayer.com/datacenters.

In this white paper, we will discuss the business and technical reasons to choose a specific data center, or at least narrow down the list to a few that are all equally suitable.

Location Relative to Users

Like many groups of factors, the reasons in this list fall into both business and technical categories. One of the advantages of the SoftLayer cloud is that resources are only placed where you choose to put them and they are never migrated by SoftLayer, giving you complete control over the location of your data.

For Latency

The first, and probably most important, consideration in selecting a data center is its location relative to those who will need to access its resources. If your employees, customers, etc. are primarily based in Canada, then a Canadian data center makes the most sense, while an Asian data center makes the least sense since the latency for everyone accessing it would be very high.

On the other hand, if you have users all over the world, you have two choices if you want to stick with a single location, namely:

- Pick where the highest concentration of users are located, which is beneficial to a large percentage of users in one geographic area. But with a small number located elsewhere, this will cause users in outlying areas to have a poorer experience.
- Select a central location where latency won't be too high for anyone but won't be very low for most users either. The idea is that few users will have a poor experience, but few will have a great experience either.

There is a third option, which is usually the best option overall: using multiple data centers located closest to your users. This provides the lowest latency, but at the highest cost due to redundant equipment in each of those data centers. Note that if you have relatively static content, such as images, video, or audio, you may consider using SoftLayer's Content Delivery Network (CDN) in conjunction with object storage. The upside of this approach is that you don't need to deploy servers with content around the world. The downside is that you will pay for every gigabyte of bandwidth used (not to mention the gigabytes used with object storage) instead of taking advantage of the free public network outgoing bandwidth that SoftLayer provides (for monthly servers only). While usually the simpler option, in extreme bandwidth consumption scenarios, this may be more expensive. You can view the locations on the CDN by going to <http://www.softlayer.com/content-delivery-network>.

You can view the real-time latencies (averaged over the last five minutes) as well as warning and critical thresholds as defined by SoftLayer based on their experience between any two SoftLayer data centers, by going to <http://lg.softlayer.com>.

For Data Privacy or Other Legal / Regulatory Reasons

The other major reason to select certain data centers over others is for various legal and regulatory reasons, such as data privacy. A common example of this is the European Union (EU) regulations around where Personally Identifiable Information (PII) can be stored for EU citizens (normally within the EU only). Russia has a similar requirement. In addition, some contracts may require data to be stored in certain countries, such as for some US federal government projects. These requirements can rapidly exclude many data centers.

Location Relative to Other Data Centers Already in Use

If you have (or plan to have) a single data center in the cloud, this section can be skipped. However, many companies have resources in multiple locations, so there are a couple of technical considerations that should be considered in selecting those data centers and/or any additional locations.

For Disaster Recovery

First, if you are deploying multiple data centers for Disaster Recovery (DR) reasons, you should consider the disasters you are trying to protect against and select locations far enough apart from each other so that a disaster that affects one location won't affect another. For example, if you are trying to protect against tornados, Dallas and Houston are far enough apart since tornados are a localized event. On the other hand, if you are trying to protect against hurricanes, they may be too close together as a hurricane that hits Houston still may cause damage in Dallas. If the problem is civil unrest / political instability, then it may not be so much a matter of distance as simply being in another country.

If you want to protect against a broad range of potential disasters, you may wish to refer to this blog post that provides some general guidance and describes the factors that should be considered:

<http://www.iso27001standard.com/blog/2012/11/19/disaster-recovery-site-what-is-the-ideal-distance-from-primary-site/>.

For Latency

Latency between the desired data centers can be a major issue when replicating data from one site to another. However, if the second data center is used solely for backup purposes or template replication, then latency isn't as much of an issue, other than potentially lengthening the backup windows.

When asynchronous replication is used, the impact of higher latency will make a slightly larger difference in the data between the two sites due to the lag between them, with a greater difference observed as latencies between the data centers increase (all else being equal). It is not a big issue if the latencies are relatively small, but when they are hundreds of milliseconds, the impact could be significantly greater, depending on the volume and rate of transactions / data to be replicated.

On the other hand, when synchronous replication is used, you may need to consider the maximum tolerable latencies supported by the applications that are doing the replication. Even if no practical limits are exceeded, the performance impact to the application will increase with each additional millisecond of latency. The reason for this is transactions need to be committed on and acknowledged from the DR site before they can be completed on the production site. This, in turn, allows for the next transaction replication to begin.

For Access to Your Existing Data Center(s)

If you have existing data centers and want to link them for data transfer purposes, you'll need to consider the same factors listed in this section in determining what SoftLayer data center(s) to use. This applies if you need to simply use data from both your own and one or more SoftLayer data centers as well as if you want to use them as part of your hybrid cloud strategy.

The following are three cloud strategies:

- Public, where all your resources are in the cloud. In this case, this section doesn't apply.
- Private, where your resources are typically on premises. In this case, this section doesn't apply to you either. On the other hand, if you use the SoftLayer definition of a Private Cloud, this may apply to you. Their definition is having your resources, dedicated to you, but hosted in SoftLayer data center(s). Many others would consider this a public cloud strategy. If all your resources are still at SoftLayer, even under their definition of a private cloud, then again, this section doesn't apply to you.

- Hybrid, where some of your resources are in your own, on-premises data centers and some are in a cloud provider's data centers. In this case, data transfer and integration of resources between the locations will matter, and you will need to consider the same issues previously described when selecting data center locations. Other issues to be considered in this model are management and monitoring options, but they don't affect the selection of data centers, and thus are outside the scope of this white paper.

Availability of Desired Features Not at Every Data Center

Another factor that must be considered in selecting a data center is the availability of various features. While most features are available at every data center, some are not, either because they are too new and are being rolled out over time (or retrofitted to existing data centers) or because they are legacy features that are being replaced with newer options (thus, they are only available in the older data centers until they are finally retired and/or the capabilities matched on the new options being deployed and those deployments completed). This section will look at options that fall into both categories.

10 Gb Ethernet

The first option is 10 gigabit Ethernet (10 GbE). All of the original data centers were connected with 1 GbE links between all the components, but the following factors have increased demand for greater throughput:

- As data volumes increase, so does the bandwidth required in order to keep down response times.
- The need for higher performance shared storage connected to more devices has required greater bandwidth on the storage devices themselves, and in many cases to the servers or VMs that they are connected to as well.
- Virtualization has led to running many VMs on a single server, especially when the virtualization platform is installed by the customer on bare metal servers. Bandwidth demands have increased in this regard in the following areas:
 - Storage bandwidth for the VMs, especially those that are storage intensive, such as database or file servers
 - Network bandwidth for the VMs, especially those that are network intensive, such as web servers
 - Access to virtualization features that recommend or require 10 GbE, such as vMotion or SMP fault tolerance on the vSphere platform
 - Use of 10 GbE on premises in customers' existing data centers, which has caused them to expect the same capabilities in their cloud providers

Not all data centers support 10 GbE yet. New data centers are being deployed with it, but some older locations have not been retrofitted with it yet. The simplest way to check on the availability of this feature is to open a ticket or use chat to query SoftLayer sales.

Auto Scale

The second area that is being rolled out, but hasn't been finished at all data centers yet, is Auto Scale. This feature will automatically create new servers (VMs) on demand and remove the servers it created if they are no longer needed, based on either resource consumption or scheduling. Currently, this feature only fully works with public cloud (multi-tenant) VMs and the SoftLayer Local Load Balancer. It can be used via APIs and/or scripting with other options as well.

This feature breaks down the ability to scale into zones and you can deploy VMs anywhere in the region using this feature, but not all data centers have this capability. For example, it can cause issues when VMs are deployed at a different data center than the rest of your equipment and there isn't a firewall deployed to protect those VMs. Another potential problem is network connectivity between data centers when Vyatta routers have been deployed to ensure that the auto scaled VMs can be accessed by the rest of your environment. They should be

protected as needed by firewalls like the existing VMs you've deployed, and load balanced across in conjunction with your existing servers.

As with 10 GbE, the simplest way to check on the availability of this feature is to open a ticket or use chat to query SoftLayer sales.

Legacy iSCSI and NAS/FTP

As some features become outdated, they are being replaced by new offerings. For instance, Legacy iSCSI has some capabilities that Consistent Performance Block storage currently doesn't have, such as storage-based replication between iSCSI devices and the ability to take snapshots. If you need these features, here are two options:

- Use legacy iSCSI. Find a data center that offers it and still has capacity.
- Use Endurance Storage from SoftLayer. It combines many of the performance options of Consistent Performance Block storage with the snapshot and replication options of legacy iSCSI. Note that this is provisioned with a defined number of IOPS per GB of space provisioned, not a fixed number of IOPS and a fixed amount of provisioned space like Consistent Performance storage uses. This is a new offering from SoftLayer and is not available at all data centers yet. You'll need to check with SoftLayer via ticket or chat as previously mentioned.

If you can't find one conveniently located that offers either the old Legacy iSCSI or the new Endurance Storage, you will need to do one of the following:

- Incur greater latency by accessing the storage across data centers (with the increase in latency that requires, it's not available with Endurance storage).
- Move/change your plans for a data center to one that has the required capabilities and capacity, subject to the other discussions on selecting a data center that will meet your needs.
- Provide the needed capability at the software level in your deployed application(s), such as by using SQL database mirroring, AlwaysOn Availability Groups, or Exchange Database Availability Groups (DAGs) to handle the replication, making use of periodic incremental backups to handle the snapshot needs, etc.

The Legacy iSCSI offering has the opposite issue of the ones previously discussed, in that the list is shrinking (in the future, the feature will no longer be available), rather than growing to eventually encompass all the data centers. Endurance block storage, on the other hand, is expanding like the other offerings described in this section. Open a ticket or use chat to query SoftLayer sales to check availability.

As with Legacy iSCSI, the NAS/FTP feature is being replaced with Consistent Performance File storage and Endurance File storage, but the new features only support NFS (commonly used with Linux), not CIFS (commonly used on Windows servers) or FTP. The solutions are similar to those with Legacy iSCSI, except you will need to deploy a file server and/or FTP server, based on protocol needs instead of accessing the device directly. These servers may be deployed using either Consistent Performance Block or Endurance Block storage if a consistent level of performance is required. Consistent File storage can't be mapped to a Windows server using NFS support, while Endurance File can be if desired. These options compare with just pointing to a NAS share maintained by SoftLayer, where no server is required for CIFS when using NAS/FTP.

Cost

Finally, all else being equal, a business driver may be the cost of the data center. The cost of every server needs to cover the operational costs of the data center, including the building, power and cooling, network bandwidth, staffing, and the actual equipment itself. Rather than having different prices at each data center, a standard price is used with a surcharge for those data centers that have higher costs associated with them in one or more of the areas previously mentioned. The surcharge can vary from as little as three percent in Montreal to as much as twenty percent in Australia.

To determine the surcharge at a data center, order a new server (it doesn't matter the type nor if it is hourly or monthly, though a few locations have monthly charges that are flat rate per server surcharges instead of a percentage surcharge), and when you select a data center from the dropdown list in the portal, the surcharge (if any) will be listed.

Conclusion

As described in this document, a range of factors can influence the data center you choose to deploy your servers in. You should consider all of the factors listed here before deciding where to place your servers to minimize the impact on performance, maximize disaster recovery capabilities, and/or the need to change locations later.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

[SoftLayer Fundamentals](#)

[SoftLayer Solutions Design](#)

[Cloud Essentials](#)

Visit www.globalknowledge.com or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

John Hales, VCP, VCP-DT, VCAP-DCA, VCI, is a VMware, SDN, and SoftLayer instructor at Global Knowledge. He is a lead instructor for the SoftLayer offerings worldwide. John is also the author of many books, from involved technical books from Sybex to exam preparation books, to many quick reference guides from BarCharts, in addition to custom courseware for individual customers. John has various certifications, including the VMware VCA-DCV, VCA-DT, VCA-Cloud, VCP, VCP-DT, VCAP-DCA, VCI, and VCI Level 2, the Microsoft MCSE, MCDBA, MOUS, and MCT, the EMC EMCSA (EMC Storage Administrator), and the CompTIA A+, Network+, and CTT+. John lives with his wife and children in Sunrise, Florida.