

WHITE PAPER: **BI IN THE CLOUD: SERVER ADMINISTRATION PRIMER**

How is server administration different in the Cloud?

Many of the challenges of maintaining a Cloud instance are the same as with a traditional server. This article will provide the details of the differences between traditional server administration and maintaining a Cloud environment. Included are tools and tips in order to deliver the strength of the Cloud using agile and scalable approaches. By better understanding the intricacies of the Cloud, you will be better positioned to leverage the strength of the cloud and avoid trouble spots.



WHAT IS THE CLOUD?

Utilizing a Cloud resource or instance is the ability to “borrow” computer power and space through the internet. Much like renting a car, you pay when you run your computer instance in the Cloud. You rent compute power from a company called a Cloud provider. As a result, you do not pay for your traditional server room and contents that whirl out hot air on weekends and holidays. With Cloud computing, less technical expertise is needed because the server instances are already configured.

The Cloud utilizes virtualization, which is a framework or methodology of dividing the resources of computer hardware into multiple execution environments, by applying one or more concepts or technologies. More simply, it is a layer of abstraction that allows the partitioning of server resources into multiple computing entities.

Cloud providers utilize software that allows for the dynamic provisioning of virtualized resources. Examples of these types of providers include Amazon Elastic Compute Cloud (Amazon EC2), RackSpace, and GoGrid. This article often uses Amazon EC2 terminology in its examples. When choosing a provider it is important to fully understand the pricing structure and the provider support history.



DISPARITY: DIVING INTO THE DIFFERENCES

How is a server started when you cannot see a power on button?

A Cloud server is started by your provider's administrative site which allows only secure users to start a virtual server with secure parameters. For example, by installing Amazon EC2 commands, a remote server can launch your Cloud instances at pre-set times such as in the morning before staff arrive. This can reduce your uptime costs by only having the server run when needed. Your server can also tell itself to shutdown at night when staff leave with the Linux cron or Windows Server at commands. If a Cloud instance is only used for a remote job, such as database loading, the startup script can include the jobs to start. Jobs can also include terminate with a shutdown after a successful batch completion.

How do you add storage when you cannot touch a hard drive?

Cloud providers allows for increasing storage space without the need to connect any hardware or bring your server down. Through your provider's interface, a virtual disk can be added that uses the traditional soft connections to connect to your instance, e.g. diskmgmt.msc for Windows and mount on Linux. While the physical hardware goes away with the Cloud services, administration of the operating system, applications, and backups still need to be done. Knowing the differences between how you support your traditional system and how the Cloud is architected can allow you to use your traditional administration knowledge with practical solutions to a new computing resource world.

For example, once the virtual disk storage is added, backups need to be done for any added data. The backups would need to move through the internet to a location where physical backups could be made, if needed. The software Rsync is very useful for this because it only moves the changes to the files which keep the network traffic low and disk storage minimal. Rsync command is native to most Linux flavors and available on Windows Server with the open source cygwin suite.

How do you install software without being able to read a CD?

Since there is no place for media, every interaction with the system needs to be done with a connection from the Cloud to a local machine with the media. This works for many applications which can be installed using web services, ftp, or ssh connections but other applications do not have this flexibility. By adding application installation as a checkpoint before moving into the Cloud, you will help assure success for your project. With this requirement you are guaranteed network accessibility to your applications that can be allow employees working options that include cell phones and laptops.

What happens when your running instance in the Cloud terminates?

Some of your data may vanish the next time you start your Cloud instance. For example, you may find that the time zone is not current and the user you added the last time the image was running is gone. Depending on your provider, some storage space is ephemeral which will not be persistent between startups.

With Amazon EC2 you have the option of rebuilding or re-bundling your image with your ephemeral changes which are often applications and operating system customizations. Once this is done, a new image can be started with the changes in place. Re-bundling can also be used for backing up the image. However, re-bundling requires that your instance be taken offline, so it is often not practical for intraday backups. You need to understand your applications and operating system changes in order to know when to re-bundle and what files to place on permanent storage. Re-bundling does have its advantages, such as with software changes. Since a reboot will erase all of your changes, if you run into an issue with a configuration change or an upgrade, you can simply terminate and restart the image to return to the previous state. Re-bundling is also a method to employ change management of your image configurations.



As simple as pressing a button can launch a development clone of a production system. This creates a low risk upgrade test platform for your applications. Many Cloud providers also do allow for launching of your instance with increased memory and processor power without any additional installation or configuration changes within your instance.

Multiple images can run at the same time as different virtual servers and become hardware clones. Your applications can be tested or upgraded without the traditional server capital and administrative setup expenses. As simple as pressing a button can launch a development clone of a production system. This creates a low risk upgrade test platform for your applications. Many Cloud providers also do allow for launching of your instance with increased memory and processor power without any additional installation or configuration changes within your instance.

Table 1: Common administrative job differences between traditional and Cloud computers.

TASK	“HOW TO” WITH A TRADITIONAL SERVER	“HOW TO” WITH A VIRTUAL CLOUD SERVER
Estimate ownership cost	Add hardware, facilities, networking, and administrator costs	Add time running with storage, memory, data transfer, and processor costs
Startup server	Push a Button	Launch from a configured web service
Shutdown	Run system shutdown command	Run system shutdown command
Change memory	Procure memory and install	Web enable increase or decrease memory
Change processors	Procure processor and install	Web enable increase or decrease processor size and speed
Change amount of storage space	Procure disk drive and install	Web enable increase or decrease storage space
Read or write to media (e.g. memory stick, cd, tape)	Place media in drive and read or write	Place media in local drive and read or write to remote Cloud through web (e.g. ftp, ssh, email)
Reproduce production system	Procure, purchase and install new hardware and software	Launch production clone with snapshot of data



AUTONOMY: SOARING WITH EASE

How can the Cloud be handled so that changes to the standard Cloud images are preserved and shared with other team members?

To keep the tuned operating system and application configurations available, place them in a shared repository. Subversion and Git, which are open source applications, are configuration management tools that work on both Linux and windows for collaboration on sharing configuration files and also productivity Cloud tools for your special business needs.

The table below gives goals for any administration environment with Cloud implementation examples. In order to reduce the likelihood of a problem, have a repeatable process that automatically handles error conditions. The fewer tasks that need to be manually done not only increase productivity but also reduce inadvertent errors.

Table 2: Good practices with Cloud examples.

GOAL	CLOUD EXAMPLES
Accessibility	Keep track of system and application configuration files for versioning and project access using a management tool such as Subversion
Delivery	Automate common administrative tasks with Linux cron and Windows at batch files
Modularity	Place files in the same application and data directories on all instances
Recoverability	Backup with rsync in multiple locations to avoid catastrophic failures
Repeatability	Document with details any changes to an initial Cloud installation
Robustness	Script using commands that work a variety of environments (e.g. Linux: <code>rm -f olddir; mkdir -pf /usr/local/bin/ec2</code>)
Usability	Place custom tools in the same directory structure and file names across all your servers
Reproductibility	Launch production clone with snapshot of data



SUMMARY

Providing proper systems administration is a critical component to moving your hardware infrastructure to the Cloud. Since much of the administration of the infrastructure is handled by the Cloud provider, it's often easy to forget that there are still many systems administration functions that falls under the responsibility of the customer. While the level of administration required varies by environment, solid administration skills are critical to application success in any environment.

About the Authors

Ms. Abramson, is an analyst at Datasource Consulting. She is an entrepreneur, technology strategist, architect and expert in Cloud, Linux, and open source strategies. With over ten year of data warehouse experience and administration, she provides independent and innovative database, statistical, and business solutions to small and medium companies. Mr. Dine is the President and founder of Datasource Consulting, LLC. He has more than 13 years of hands-on experience delivering and managing successful, highly scalable and maintainable data integration and business intelligence solutions. Steve is a faculty member at The Data Warehouse Institute and a judge for the Annual TDWI Best Practices Awards.