Spot Pricing Analysis

Introduction

Amazon provides a paying possibility for EC2 instances called Spot Instances. This means that customers can bid on unused EC2 capacity, and that they will get a running instance as long as their bid exceeds the current spot price. Amazon claims that these spot prices change according to the rules of supply and demand, but whether this is fully the case is however questionable. If a customer has an application that is flexible enough to allow it to cease its operations during periods when the customers bid price does no longer meet or exceed the spot price of that moment. This pricing model is thus only usable for applications under certain constraints and they have to be adapted to be able to recover from instance outages (e.g. through checkpointing). Using Spot instances can however significantly lower your costs, since prices most of the time are a lot lower than On-Demand prices and even than Reserved instance prices.

Amazon released this innovative pricing (auction-like) mechanism in December 2009, so since it is still rather new not that much research has been done. To perform statistical analysis on the spot pricing, we need the history of these prices. On Amazon these prices can be acquired by the 'ec2-describe-spot-price-history' webservice call. On cloudexchange.org this data can be downloaded as a CSV which provides the full history of the spot prices (for all instances except for Micro instance). Also note that cluster compute spot instances do not exist at this time.

Spot instances open a market model where unused resources can be sold through the mechanism of supply and demand. Or not? The problem is it does not have to be supply/demand driven because there is only one seller. No one can assure us that Amazon does not use all the information it has at hand (like the maximum bids customers placed on the spot market) to determine the current spot price. Applications that need instances urgently can bid more get the remaining resources at its disposal, specifying a higher maximum bid raises the priority of a request for capacity. Why wouldn't one always bid the on-demand price on the spot market? The problem with this technique is that since the spot price can actually rise above the on-demand price (as the history shows us) all your instances can get terminated at one time. Spot Instances however are especially applicable to certain applications that are flexible (are easily stopped and started again) such as financial modeling, web crawling, load testing, video conversion, … These tasks can be performed in iterations, which makes taking snapshots easy and thus are these types of applications resilient to the fact that your instance can just be killed when the current spot price exceeds your maximum bid.

Working method

Looking at the spot price history will be done for the US East region, since most fluctuating graphs were generated for this region. The different scenarios of dividing the spot instance prices we're having a look at are (for the last 2 months):

- Average per day
- Average per week
- Average per day of the week
- Average per hour of the day
- Average for the day versus the night

After generating the required graphs (including blox pots) for these scenarios, we'll also have a look at how this data can be used to make the EC2 capacity planner more intelligent (e.g. it will foresee cheaper prices at certain times of the day). To make this possible some statistics are extracted from the spot price history data set.

The statistic values that are stored for every scenario (and this for each OS-Region-Instance type combination possible) are:

- Kurtosis: "it is a measure of the 'peakedness' of the probability distribution of a real-valued random variable. Higher kurtosis means more of the variance is the result of infrequent extreme deviations, as opposed to frequent modestly sized deviations."
- Skewness: "it is a measure of the asymmetry of the probability distribution of a real-valued random variable. The skewness value can be positive or negative, or even undefined. Qualitatively, a negative skew indicates that the tail on the left side of the probability density function is longer than the right side and the bulk of the values (including the median) lie to the right of the mean. A positive skew indicates that the tail on the right side is longer than the left side and the bulk of the values lie to the left of the mean. A zero value indicates that the values are relatively evenly distributed on both sides of the mean, typically but not necessarily implying a symmetric distribution."
- Arithmetic mean: "it is often referred to as simply the mean or average when the context is clear, is a method to derive the central tendency of a sample space."
- Geometric mean: "it indicates the central tendency or typical value of a set of numbers. It is similar to the arithmetic mean, which is what most people think of with the word "average", except that the numbers are multiplied and then the nth root (where n is the count of numbers in the set) of the resulting product is taken."
- Number of Values: "it is an indication of the number of price changes."
- Maximum: "it is the greatest value in the set."
- Minimum: "it is the smallest value in the set."
- Percentile: "it is the value of a variable below which a certain percent of observations fall. The 25th percentile is also known as the first quartile (Q1); the 50th percentile as the median or second quartile (Q2); the 75th percentile as the third quartile (Q3)."
- Variance: "it is used as one of several descriptors of a distribution. It describes how far values lie from the mean."
- Standard deviation: "it is a widely used measurement of variability or diversity used in statistics and probability theory. It shows how much variation or 'dispersion' there is from the 'average' (mean, or expected/budgeted value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data is spread out over a large range of values."

Spot Prices (US East Region)

We'll have a look at the created graphs for the us-east-1.linux.m2.xlarge instance.

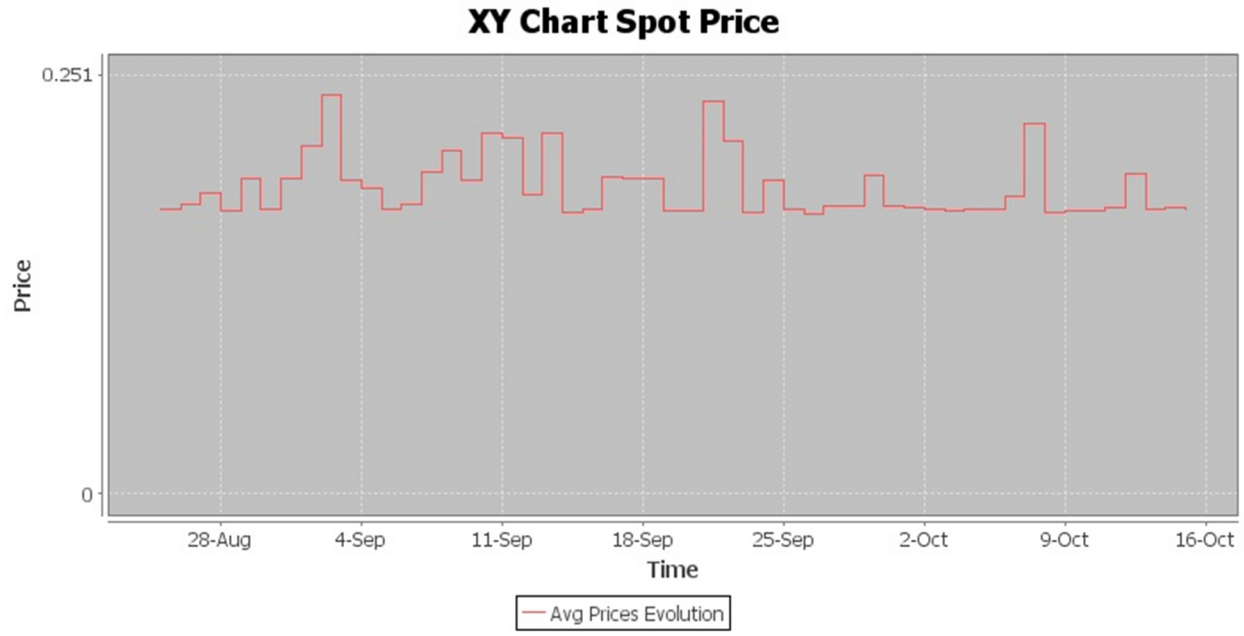

This first graph gives an overview of the last 2 months what the average price was for each date. We notice that it seems to fluctuate quite a bit. When we have a look at the box plot for this data we see that the percentiles are aligned quit well, but that the averages sometimes increase a lot because of outliers. This is also easy to tell by having a look at the acquired statistics for this instance type: Q1=0.165, Q2=0.172, Q3=0.176 but it has an arithmetic mean of 0.1833. The Kurtosis is 21.01, which tells us that there are infrequent extreme variances (outliers). We can conclude from this data that the statistics can tell us whether a current price is acceptable or probably an outlier.

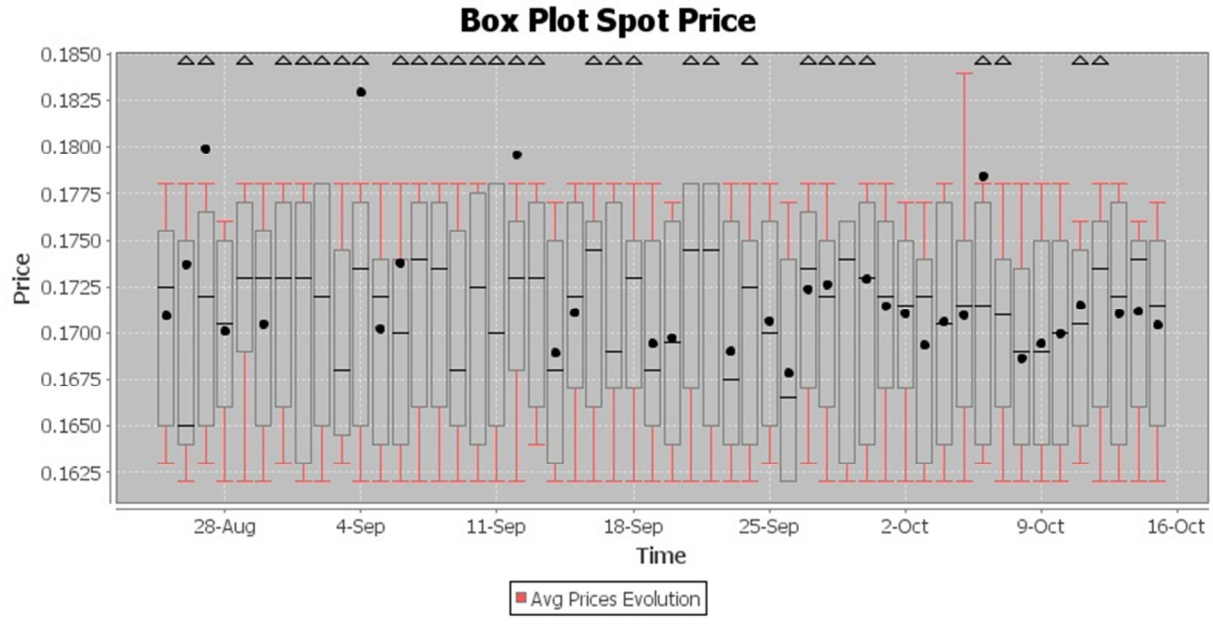

If we look at the average spot price per hour during this period, we again would think that it fluctuates a lot. Also it seems that prices are higher in the afternoon and evening, while they are cheaper during the night and the morning.

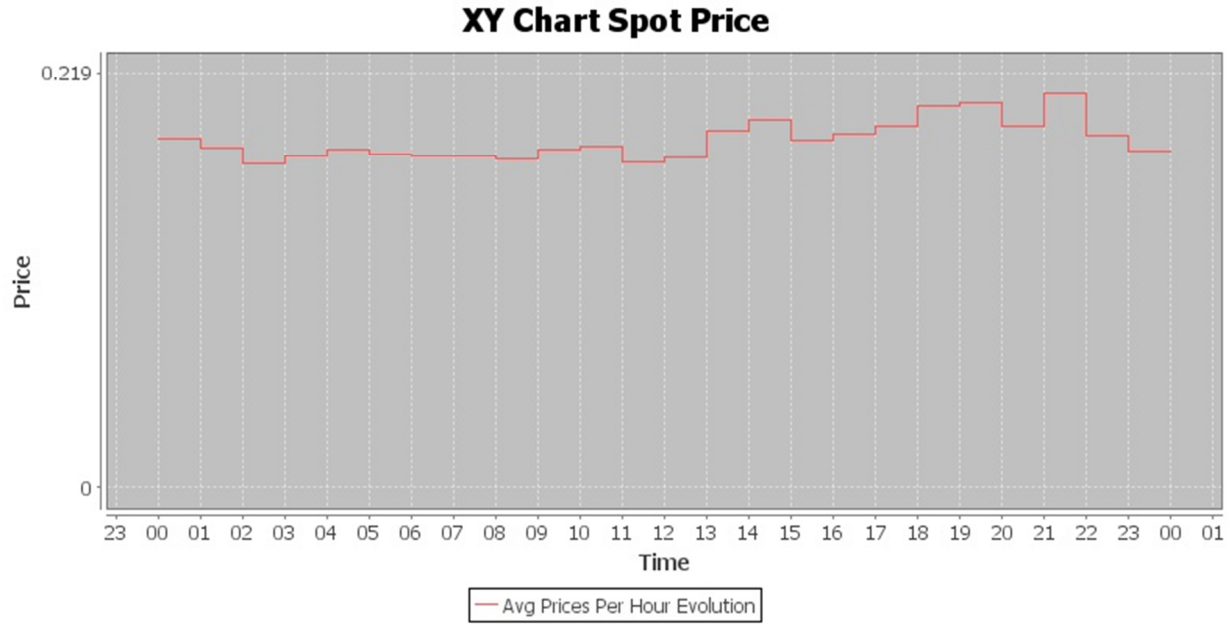

But if we look at the box plot it is again clear that the percentiles lay rather close to each other, and thus that the average spot price fluctuations are caused by outliers. The statistic values are almost always around 0.165 for Q1 and around 0.176 for Q3. The Kurtosis value is rather high most of the time, which indicates the presence of outliers.

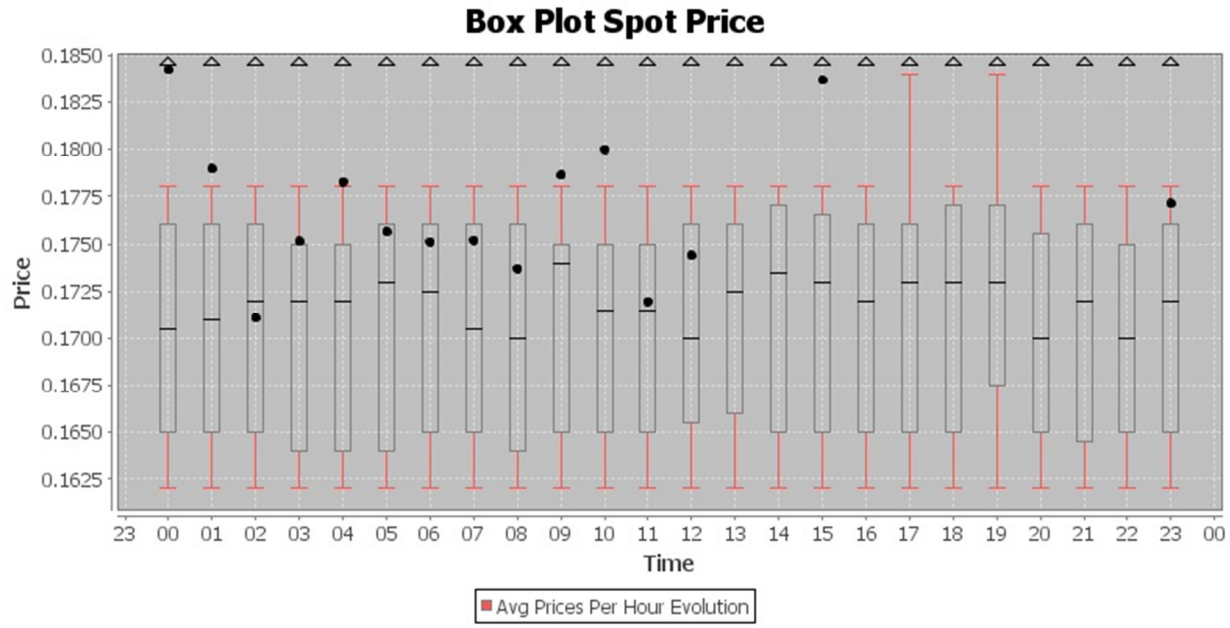

Next we have a look at how the price evolution during the days of the week.

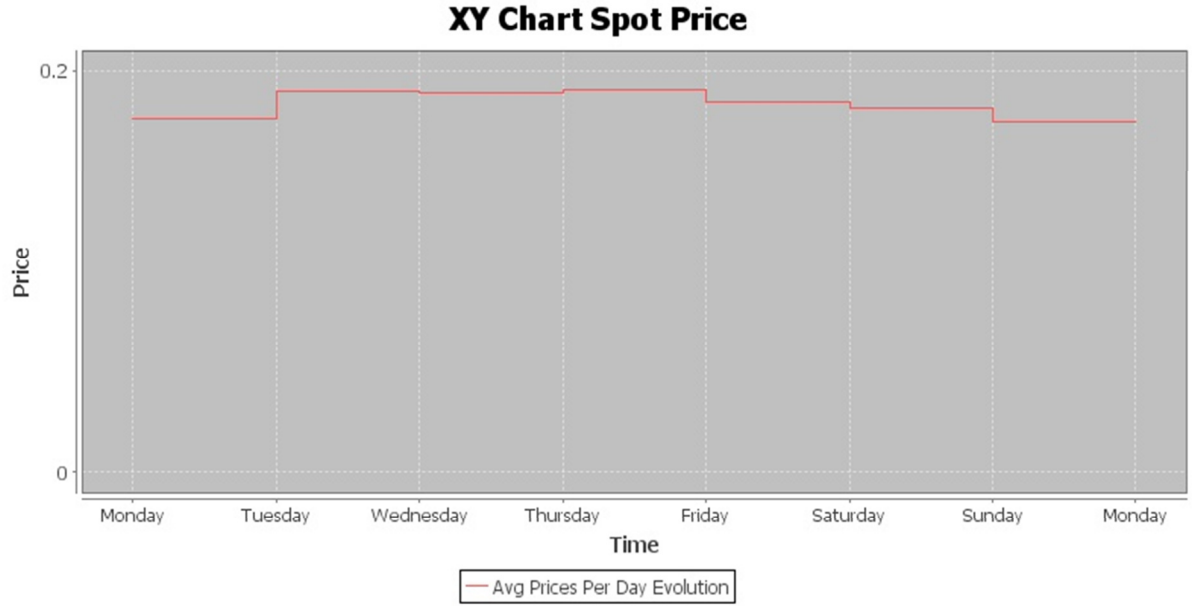

For this instance the average price is lower from Friday till Monday, but if we have a look at the boxplots below we notice that 75% (Q3) of all values fall below 0.175 during the weekend while this border is higher during the week. So we can conclude that it would be intelligent to assume weekends to be cheaper than weekdays for this instance, at least when an extra check is performed to see whether a current price is an outlier or not. The statistic values confirm this observation: the Q3 value for the weekend days is lower than the same value for the week days. And for all days a high Kurtosis value is obtained, which indicates the presence of outliers. Note that this fact is also observable on the previous graph (average spot price per date), we notice that there never is a peak in average price in the weekends (since weekends start at the vertical grid lines).

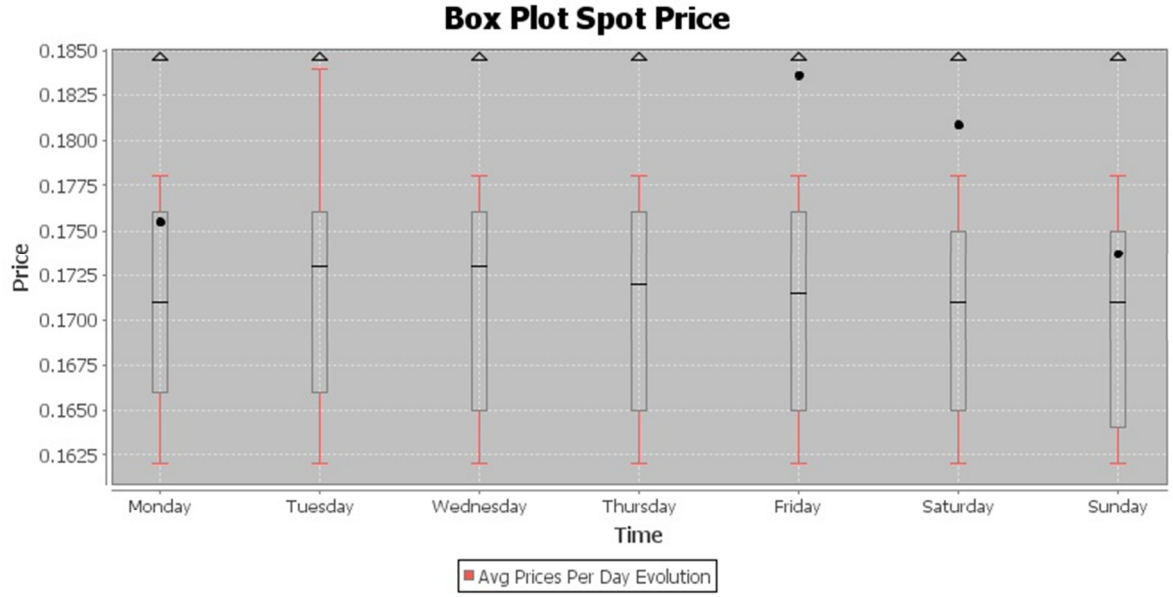

If a difference in price between the weekend and week days is noticeable then it is rather evident to have a look at the price difference between day and night:

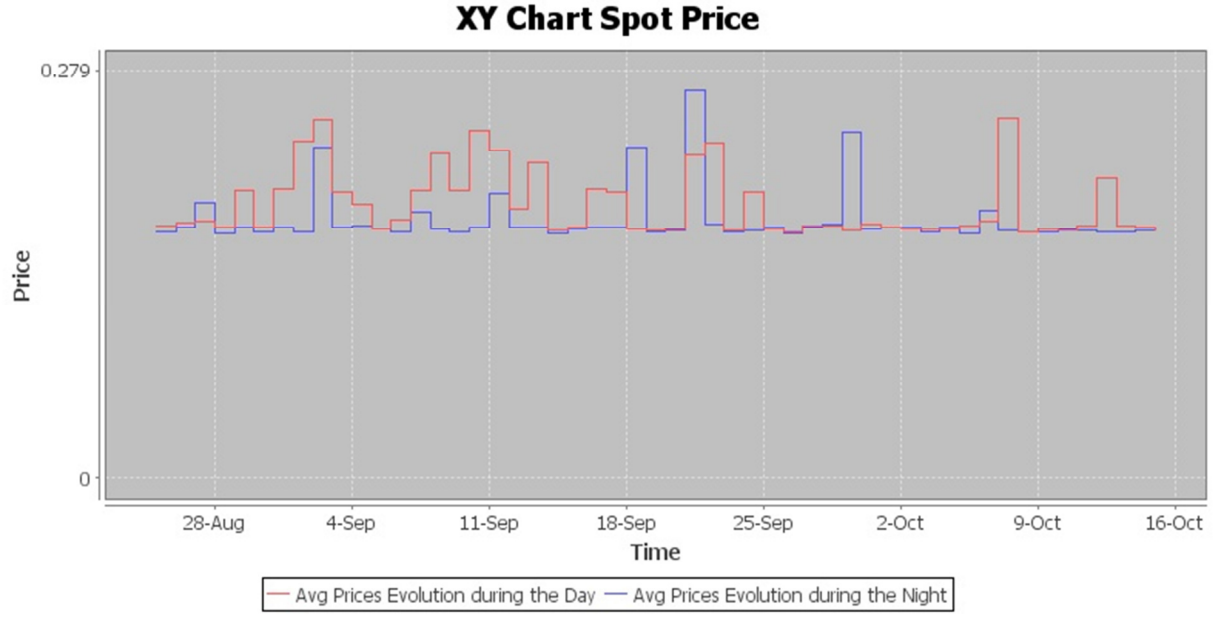

The plot beneath shows the difference between the average spot price during the day and the average spot price during the night for that day. A positive peak means that the spot price was more expensive during the day. Note that this was the case for most of the days. If we have a look at the statistics we see that the average difference is 0.0096 (which is rather small if the total average equals 0.1833) So we conclude that our suspicion of higher prices during the day is confirmed, but that the difference is very small. But it is probably not a bad idea to bear in mind that this observation could become more clear in the future, when the spot market becomes more active. The average difference and the standard deviation are probably a good metric to tell how much day versus night prices fluctuate.

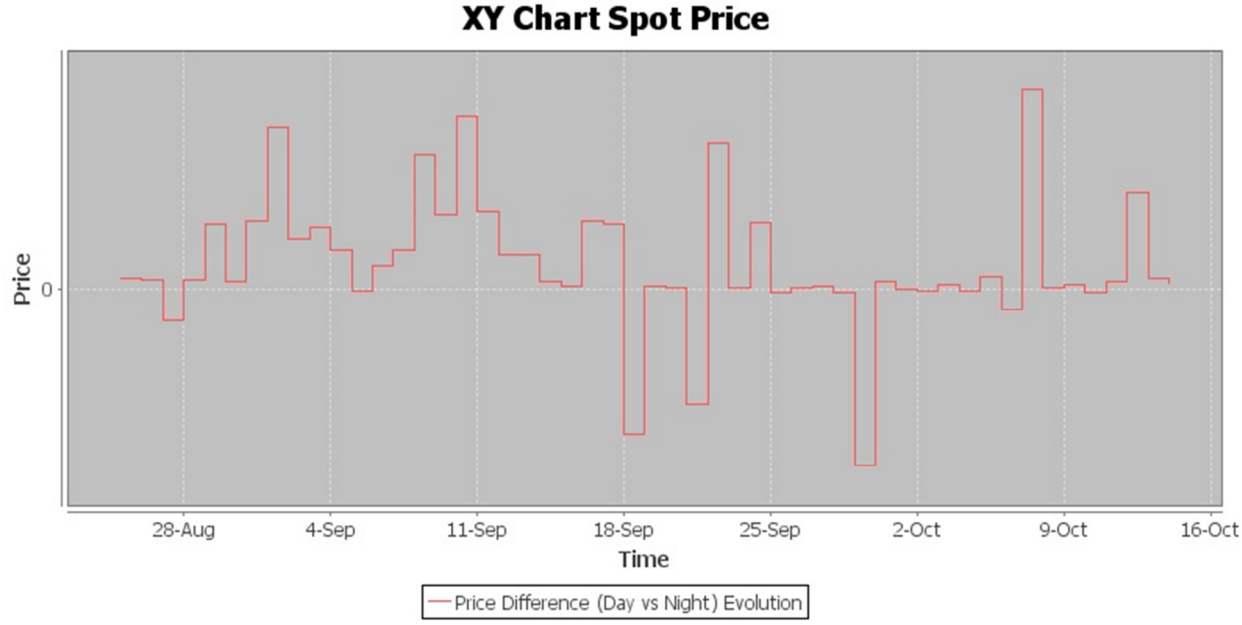

Now we'll have a look at the average spot price per week graph, again we see fluctuations. Now we have seen that these fluctuation are present when we look per hour, per day and per week already. So the market mechanisms are active for this instance for some time already and on every applicable time scale.

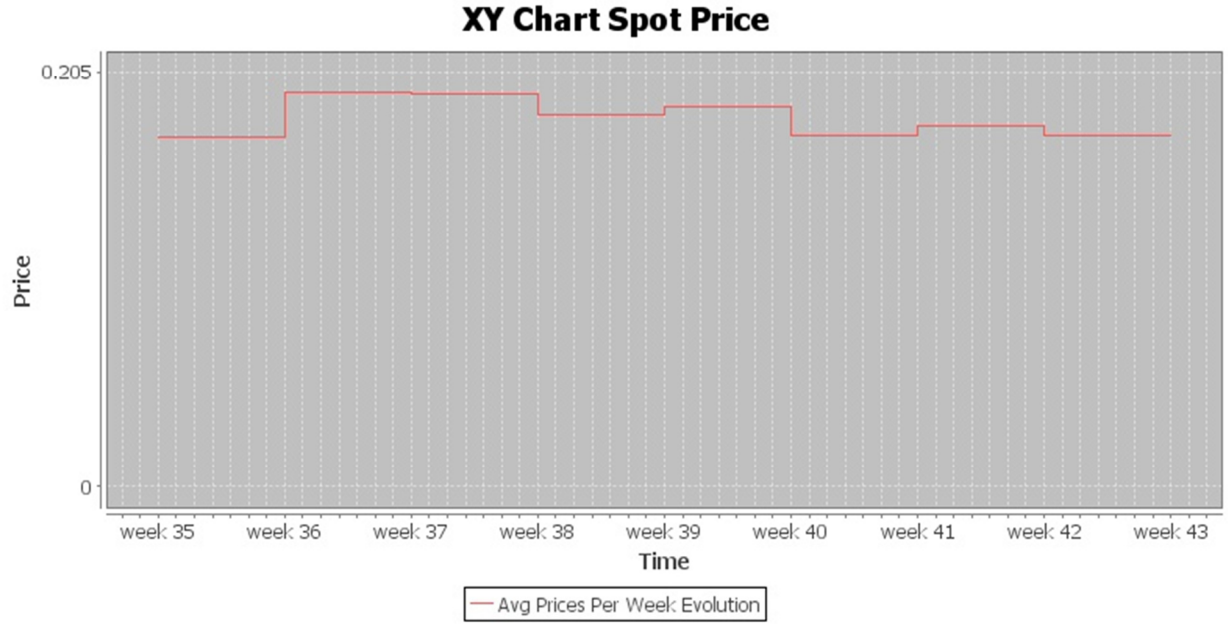

Once again the boxplot graph shows less fluctuations. These differences in average spot price are again caused by outliers (high Kurtosis values).

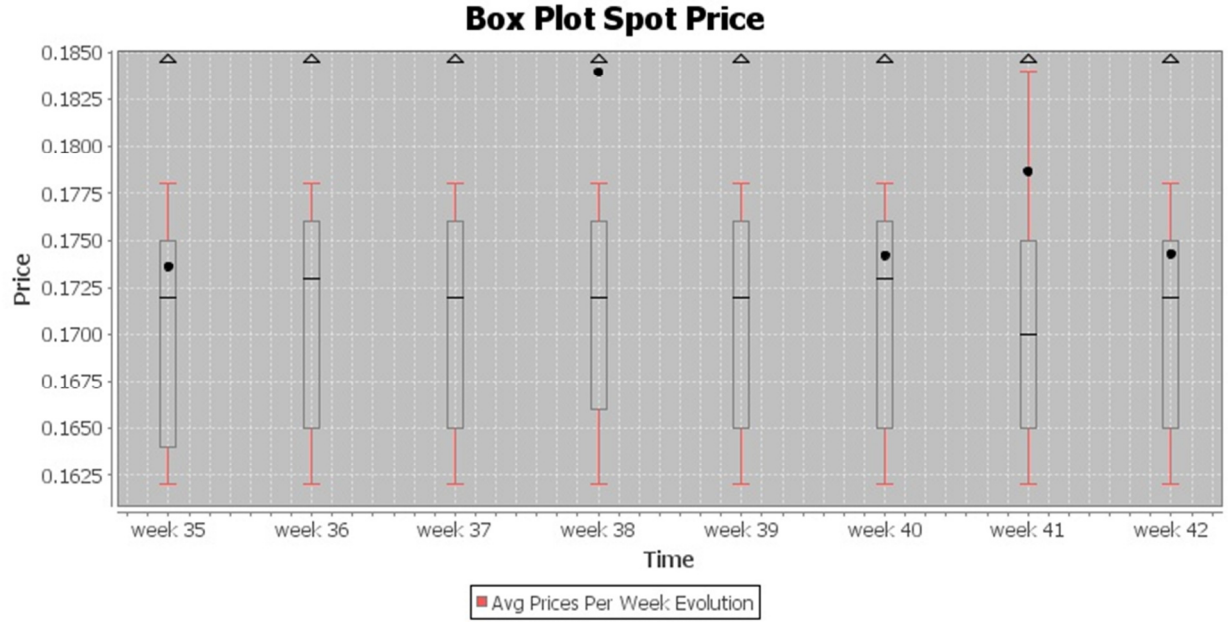

Now we'll compare these results with other instance type prices in the US East region. The following graphs were created for the us-east-1.linux.c1.medium instance type. We notice that the average stays rather constant (flat graph) till the beginning of October. This probably indicates less activity on the spot price market before the beginning of October. The box plot grapth confirms that it only starts to become interesting in the beginning of October, since the percentiles start to be different. A good measure of market activeness is probably the average difference in Q1/Q2/Q3 values.

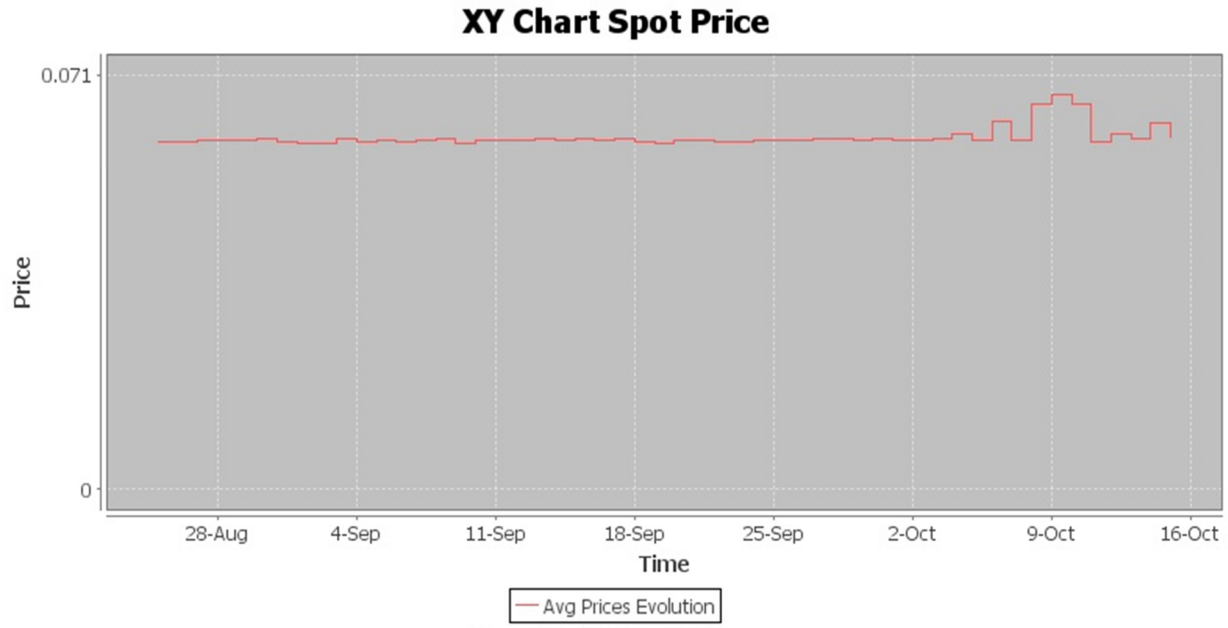

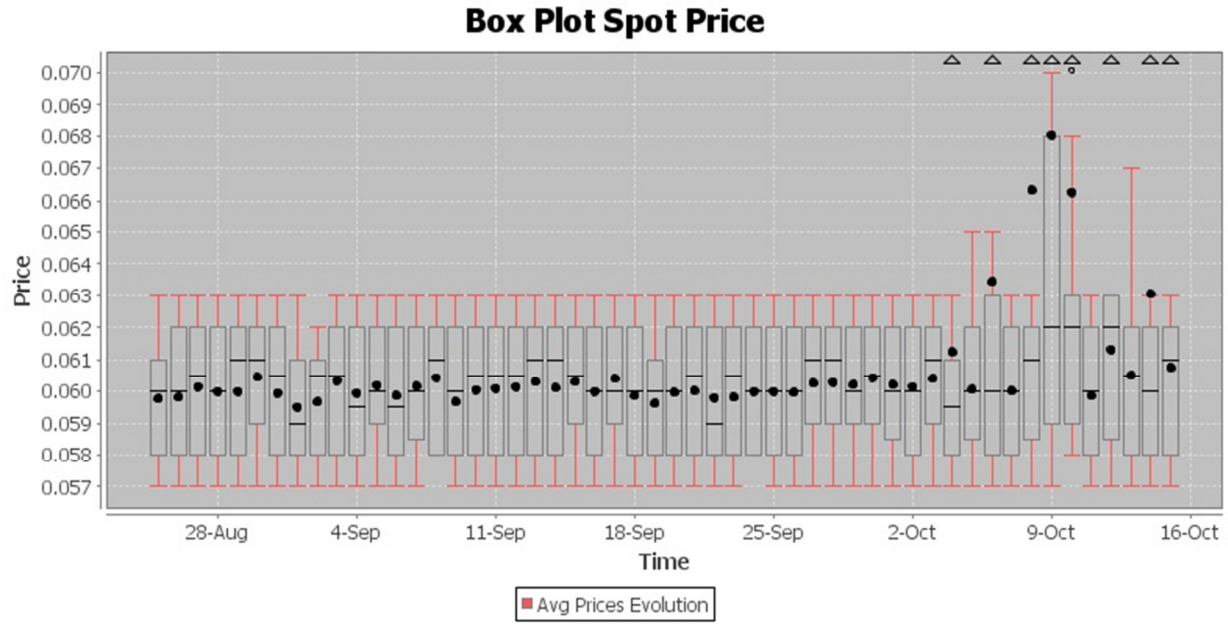

Also other graphs for this instance type start to be interesting in the beginning of October. This shows that it is important to look at data for different time frames because it could be that at certain times the market mechanisms are not fully operating yet. We notice in the statistic values that only for week 41 and 42 a high Kurtosis value is calculated, so no outliers occured before.