

Cloud Performance Benchmark Series

Amazon EC2 Web Serving

Md. Borhan Uddin

Bo He

Radu Sion

ver. 0.9c



1. Overview

Experiments were performed to stress test HTTP performance on Amazon EC2 (Elastic Compute Cloud). Testing was done in three different setups: intra-region (HTTP server and probing machine in the same EC2 region), inter-region (HTTP server and probing machine in different EC2 regions) and probing from local wired network. HTTP servers were running on small, medium and large instances of EC2 machines in all four Amazon EC2 regions (US East, US West, EU and Asia).

2. Setup

LAMP (Linux, Apache, MySQL, PHP) HTTP servers running on small, medium and large instances of EC2 machines were probed from another large instance of EC2 machine with JMeter (jakarta-jmeter-2.3.4) installed on it for inter and intra region setup; and from local Linux/Ubuntu boxes.

At the time of testing, the specifications of the Small, Medium and Large Amazon Machine Images (AMIs) were as follows:

Small (Default)

LAMP Web Starter (AMI Id: ami-2cb05345)

Fedora Core 8, 32-bit architecture, PHP 5.2, Apache 2.2, and MySQL 5

Small (m1.small) CPU Units : 1 ECU, CPU Cores :1 Core, Memory: 1.7 GB

Medium

LAMP Web Starter (AMI Id: ami-2cb05345)

Fedora Core 8, 32-bit architecture, PHP 5.2, Apache 2.2, and MySQL 5

High CPU-Medium (c1.medium) CPU Units: 5 ECU, CPU Cores: 2 Cores, Memory: 1.7GB)

Large

Basic 64-bit Fedora Core 8 (AMI Id: ami-86db39ef) with LAMP installed on it.

Fedora Core 8, 64-bit architecture, and Amazon EC2 AMI tools.

Large (m1.large) CPU Units (4 ECUs), CPU Cores (2 Cores), Memory (7.5 GB)

The local probing machines' configurations were as follows:

2.x GHz Intel Core 2 Duo processors, 2-4 GB RAM, 100/1000 Mbps Ethernet LAN, Redhat Fedora 12 / Ubuntu 9.04, 32-bit architectures with Sun Java VM 1.5+, Jakarta JMeter 2.3.4+ and PERL installed.

2.1. Throughput

In the probing machines (large AMI with above configuration for intra and inter-cloud setup, local Ubuntu setup), JMeter was installed together with a custom PERL for HTTP probing of servers running on small, medium and large instances. The PERL script invokes the non-GUI mode of JMeter for testing different configurations and loads on HTTP server. The number of concurrent users and number of requests per user are varied to generate different HTTP server loads. The number of concurrent users is varied from 1 to $10^{\text{MAX_TH}}$; the number of requests per user are varied from $10^{\text{MAX_TH}}$ to 1. MAX_TH is adjusted to match the probed instance types and type of setup as inter, intra-cloud and local setup (for a typical setup MAX_TH was between 10 and 14). Accordingly, the number of concurrent users was varied from 1 to 1024-16384 and number of requests per users were varied from 1024-16384 to 1 based on the types of setups and instances.

Trace-routing of the hosts, running the JMeter for different loads, result files backup and summary extractions are performed by the custom PERL script. The HTTP server load is varied from very low to very high. For a given configuration, its maximum seen throughput (connections/sec), average connection size and achieved bandwidth (in KB/sec and in Mbps) are extracted. Four different Amazon EC2 regions were targeted and three types of instances (small, medium and large) were tested in each, for a total of 12 intra-region, 12 local and $12 \times 3 = 36$ inter-region setups.

2.2. Transaction Cost

In addition to core throughput (connections/second) for each setup, the dollar cost per HTTP transaction/connection is also inferred for the Amazon pricing model at the time of the experimentation (included here for reference). Moreover, it is important to consider the 3 types of EC2 instance service levels with different pricing: On-Demand, 1-year Reserved, 3-year Reserved. For running HTTP server on Amazon, pricing of two types services need to be considered- instance cost and cost of data transfer –in and out. The pricing of EC2 instances & data transfer cost were collected from Amazon ([HTTP://aws.amazon.com/ec2/#pricing](http://aws.amazon.com/ec2/#pricing) and [HTTP://aws.amazon.com/calculator](http://aws.amazon.com/calculator)) for all three types of instance services. Figure 1 shows an example hourly pricing chart of Amazon EC2 On-Demand instances.

US – N. Virginia	US – N. California	EU – Ireland	APAC – Singapore
Standard On-Demand Instances		Linux/UNIX Usage	Windows Usage
Small (Default)		\$0.085 per hour	\$0.12 per hour
Large		\$0.34 per hour	\$0.48 per hour
Extra Large		\$0.68 per hour	\$0.96 per hour
High-Memory On-Demand Instances			
Extra Large		\$0.50 per hour	\$0.62 per hour
Double Extra Large		\$1.20 per hour	\$1.44 per hour
Quadruple Extra Large		\$2.40 per hour	\$2.88 per hour
High-CPU On-Demand Instances			
Medium		\$0.17 per hour	\$0.29 per hour
Extra Large		\$0.68 per hour	\$1.16 per hour

Figure 1: Pricing (Hourly Rate) of Amazon EC2 On-demand machines (table from amazon.com)

The cost per transaction/connection is calculated in US microcents as follows:

Cost per connection in On-demand instances (in microCent) =

$$\left(\frac{\text{Hourly instance rate}}{(\text{Throughput}) * 3600} + (\text{Data Transfer price rate} * \text{Amount of Data Transferred}) \right) * (10^8)$$

Cost per connection in 1 – year Reserved instances (in microCent) =

$$\left(\frac{\text{Hourly instance rate}}{(\text{Throughput}) * 3600} + (\text{Data Transfer price rate} * \text{Amount of Data Transferred}) + \frac{1\text{-year contract cost}}{\text{Conn per second} * 365 * 24 * 3600} \right) * (10^8)$$

Cost per connection in 3 – year Reserved instances (in microCent) =

$$\left(\frac{\text{Hourly instance rate}}{(\text{Throughput}) * 3600} + (\text{Data Transfer price rate} * \text{Amount of Data Transferred}) + \frac{3\text{-year contract cost}}{\text{Conn per second} * 3 * 365 * 24 * 3600} \right) * (10^8)$$

3. Results

3.1. Throughput

Intra-region. In this setup the probing machines and probed HTTP server were in the same EC2 region; the distance between probing machine and HTTP server was between 1-4 hops. Regions exhibited a throughput distributed as expected. Small instances showed lower throughputs around 275-315 conns/sec, and medium or large instances hovered around 500-600 conns/sec.

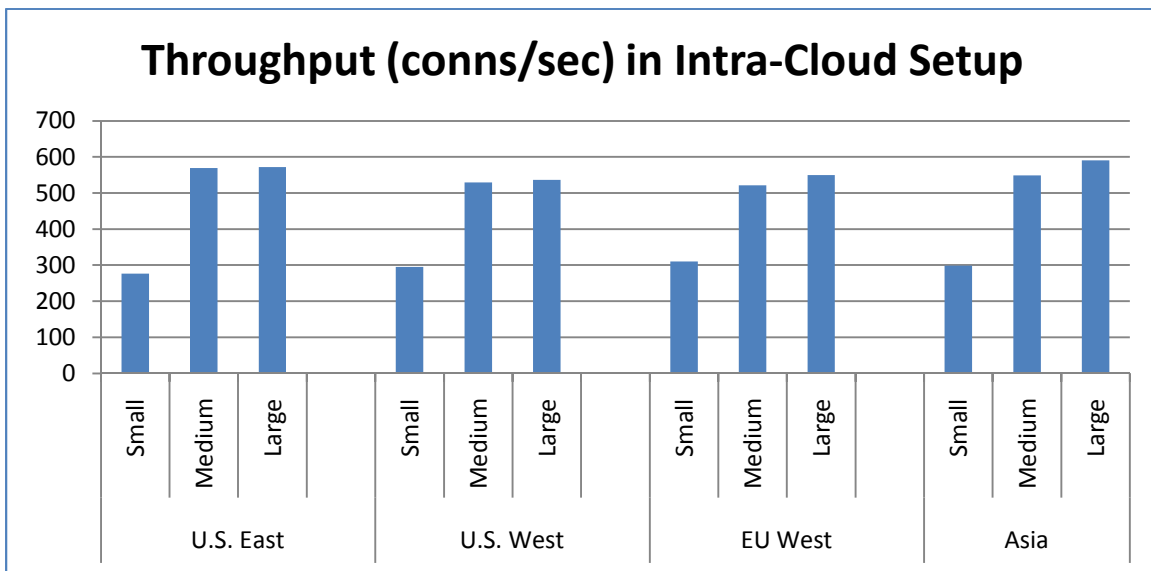


Figure 2: Throughput (Connections/second) in Intra-Cloud Setup for 4 different EC2 regions.

No significant difference in throughput between medium and large instances was observed, however. This may be due to network bandwidth shaping and/or co-hosting of medium/large instances on the same hardware. For small instances, achieved bandwidth is around 12k-13k KB/sec i.e. approximately 100Mbps and for medium or large instances achieved bandwidth is around 20k-25k KB/sec i.e., 165-200Mbps. Figure 2 depicts the throughput (connections/second) of 4 different EC2 regions for intra-region setup.

Inter-region. Notably, in this setup, throughput does **not** depend much on HTTP server configuration, but mostly on network link delays. *Sometimes*, experiments showed that **small instances show better throughput than medium or large instances** because of different link delays. I.e., small instances throughputs of 48-220 conns/sec (corresponding to a traffic bandwidth of 15-80Mbps) were achieved. Medium and large

instances reached up to 100-243 conns/sec (achieved bandwidth 30-90Mbps). Figure 3 shows the throughput (connections/second) for inter-region setup.

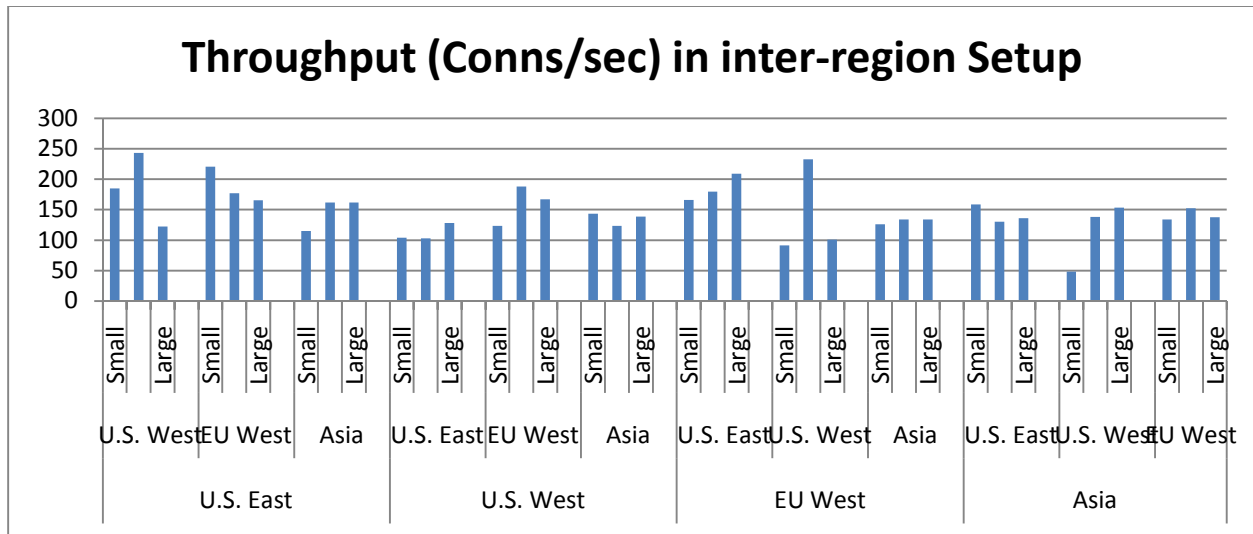


Figure 3: Throughput (connections/second) for inter-region setup. Region names at the bottom indicate location of probing machines.

Traceroute information between the local probing machines and servers was used to confirm that **routes between the regions are sometimes asymmetric** and **link delays vary significantly between the routes, and in time**. Also, throughput for same pair of regions with respect to the probing machine's location is often different, e.g., throughput of US East→EU is different from EU→US East. No clear-cut relationship between throughput and proximity of regions in the inter-region setup was found.

Local Setup. The lowest throughput (17-26 conns/sec) was achieved when probing remotely from the local setup despite it being connected directly through multi GBps low latency links to the main eastern board US ISPs. We suspect a traffic shaping router somewhere on the path. Moreover, traceroute/ICMP packets were also dropped. We will attempt to replicate these results in a different environment for confirmation. Accordingly, achieved bandwidth for these tests was 5-8 Mbps.

For the US East EC2 region (the most nearby EC2 region to the test-point) the highest throughput was observed which leads us to believe that traffic shaping / connection throttling may happen inside (or close to) EC2. The throughputs of **targeting both the US East and US West region seem consistent with the configuration of the target machines**. Larger instances score higher, and smaller instances score lower. Throughput to **EU West and Asia was not consistent with the target machine configurations** likely

due to link delays between the servers and the local probing network. Figure 4 shows the throughputs (connections/second) for this setup.

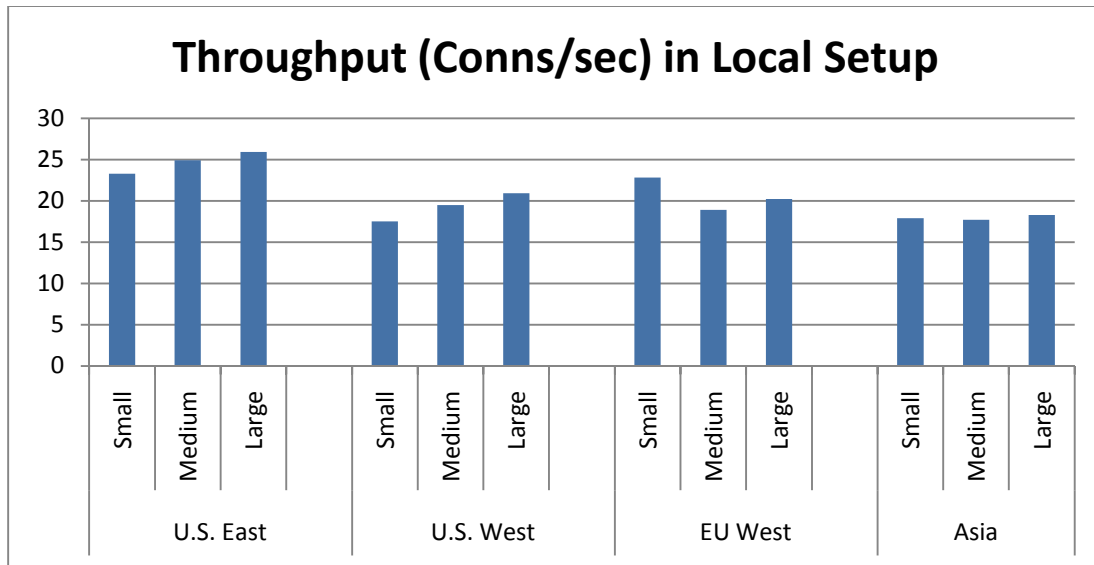


Figure 4: Throughput (connections/second) for local wired network setup for 4 different EC2 regions.

To validate this observed low throughput behavior, and eliminate any local interference, we performed the same experiments from different vantage points and different ISPs. The observed behavior was identical. Figure 5 illustrates the throughput from a high-speed ISP connected *geographically close* vantage point to the US East EC2 Region.

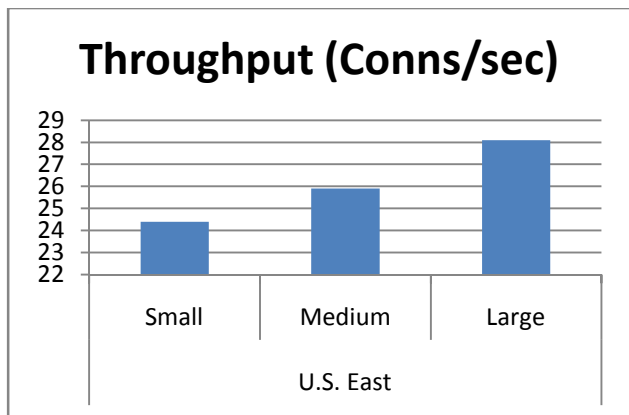


Figure 5: Throughput (connections/second) to servers hosted in the US East EC2 region from a different out-of-cloud *geographically close-by* vantage point.

3.2. Cost of Transaction/HTTP Connection

Intra-region. In this setup, large instances seem to exhibit the lowest cost/connection (385-398 microcents), less than small and medium instances (422-432 microCent). However, there seems to be *no significant difference between the cost/connection of*

small and medium instances. For all regions and all server configurations, seemingly 3-year reserved contracts offer the cheapest cost/connection. Also 1-year contracts offer a cheaper deal than On-demand instances. However, **the cost/connection differences among the three types of contract are often negligible.** Other factors need to be carefully considered before purchasing reserved instances! Figure 5 illustrates the cost/connection data.

Inter-region. In this setup costs/connection do not depend much on server configuration. Seemingly the cost/connection is lower for 3-year than for 1-year contracts. Also the cost/connection in 1-year contracts is lower than for on-demand instances, although the differences are not as high as in the intra-region setup. Notably, **the cost/connection is not symmetrical**, e.g., US East→US West and US West→US East are priced differently (as a direct result of asymmetry in throughput) etc. Geographically proximity also doesn't seem to affect the cost/connection. Figure 6 illustrates this data.

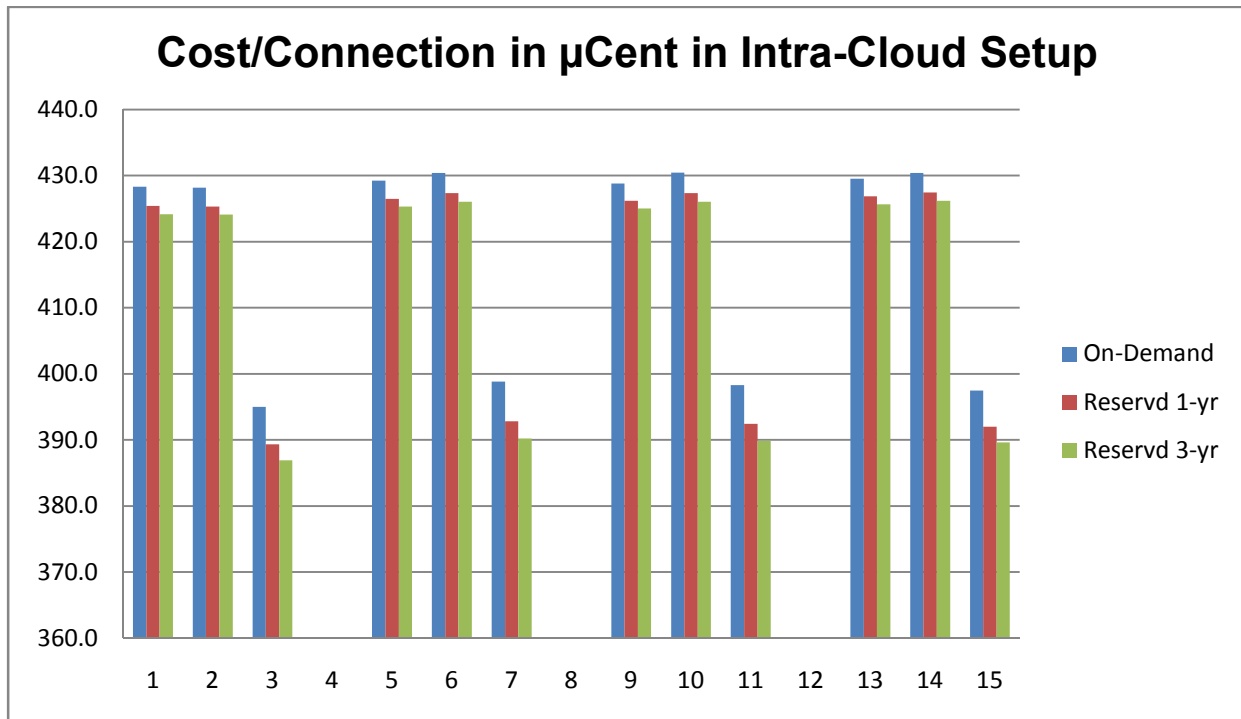


Figure 6: Cost/Connection in US microcents for intra-region setup for 4 EC2 regions.

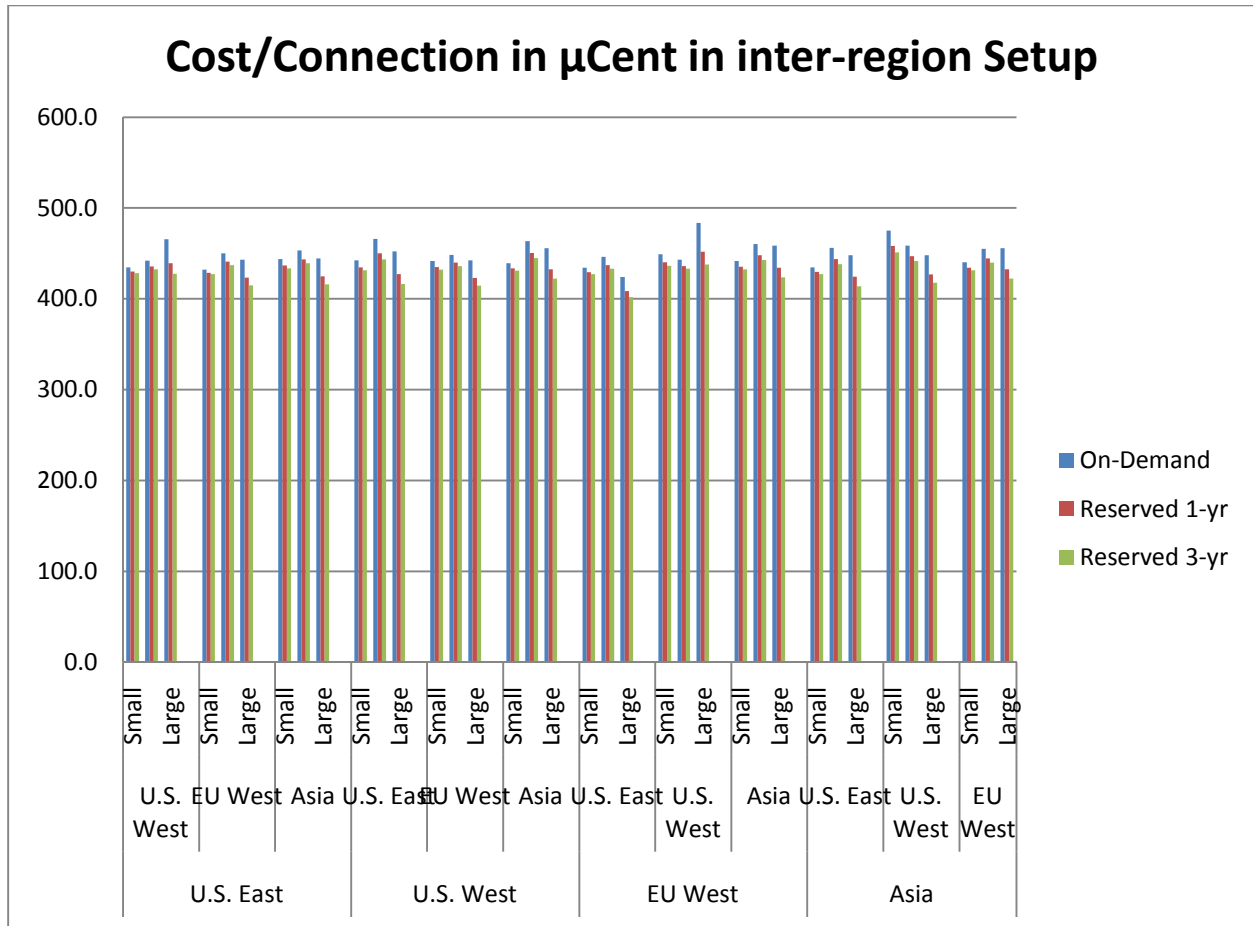


Figure 7: Cost/connection (in micro Cent) for inter-region setup. Region Names at the bottom indicate location of probing machine.

Local setup. As expected the low throughput results in a high cost/connection. Another phenomenon observed here comes from the fact that traffic shaping / connection throttling results in a low throughput that does not change for different machine configurations. As a result, per-connection, large and medium instances are costlier than smaller instances. Figure 7 illustrates.

Figures 9 and 10 illustrate the same phenomena when connecting to the US East EC2 region, from a different local vantage point.

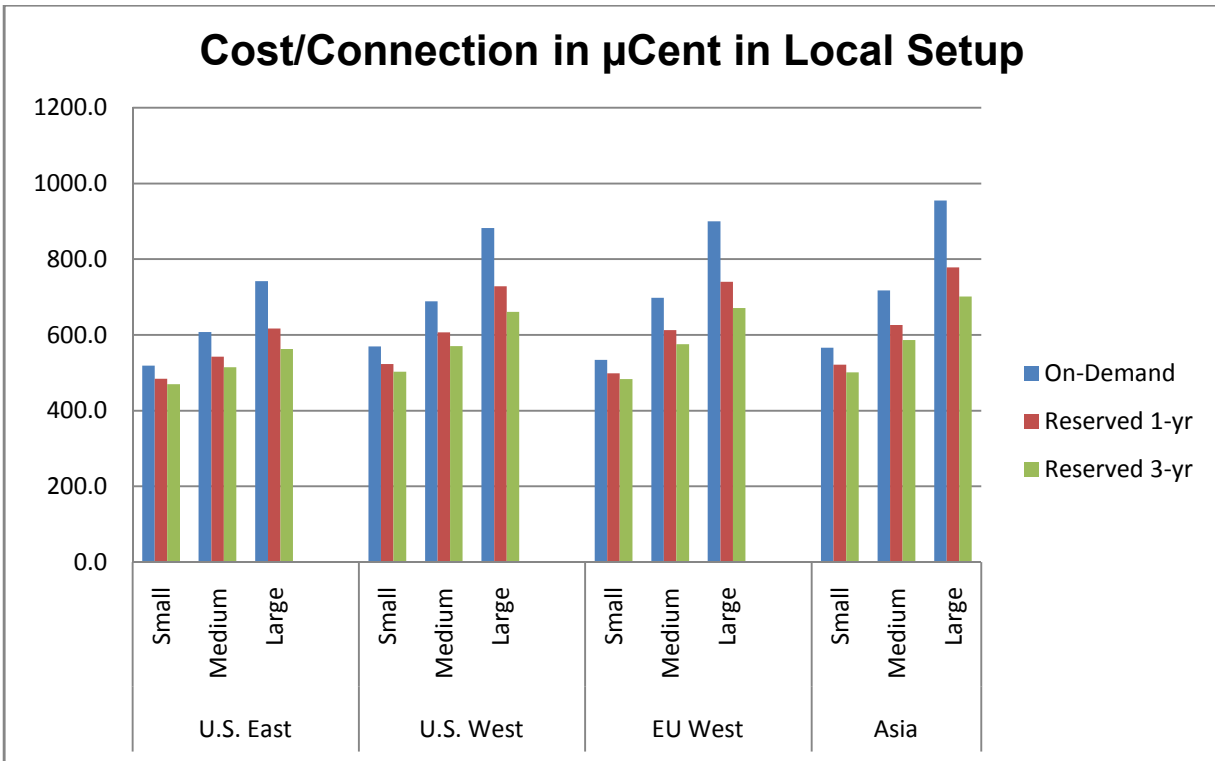


Figure 8: Cost/Connection in US microcents for local setup probing 4 EC2 regions.

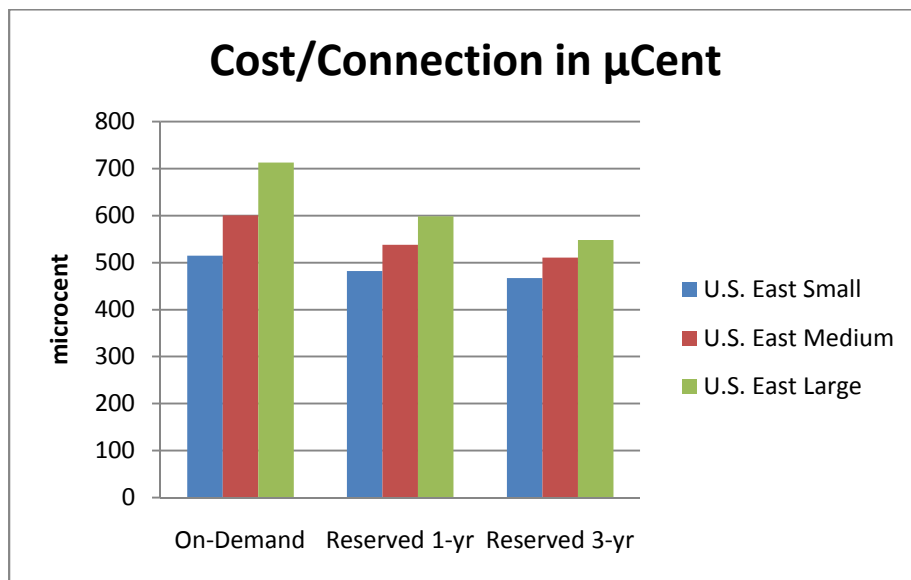


Figure 9: Cost/Connection in US microcents for (second) east coast local setup testing point probing East Coast EC2 regions.

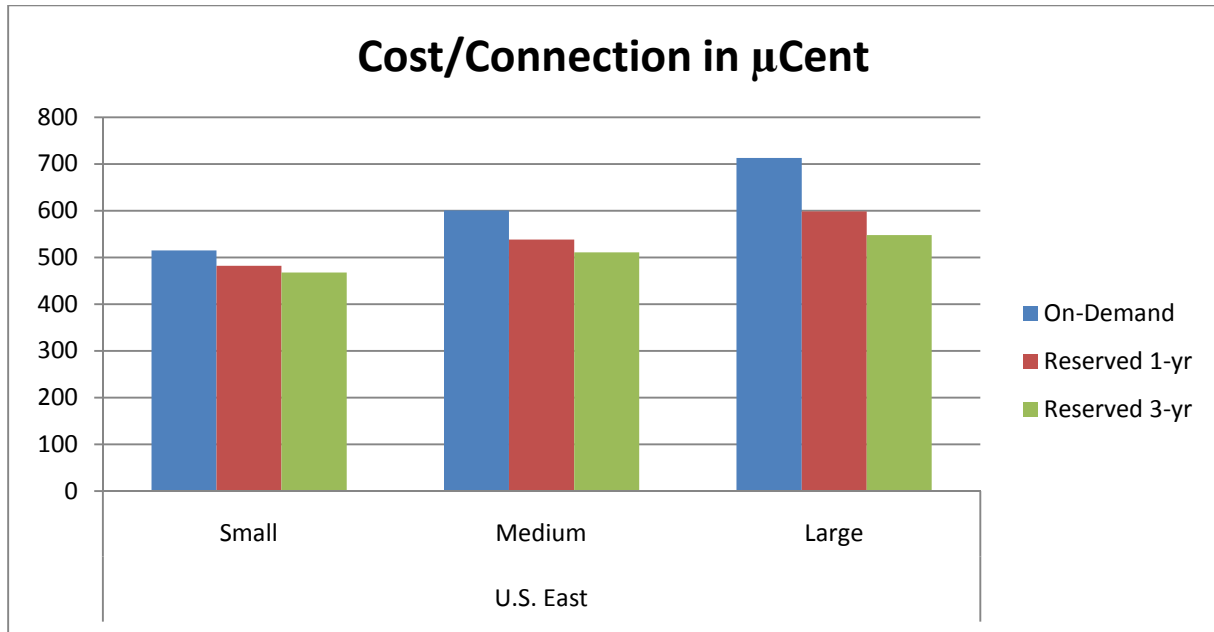


Figure 10: Cost/Connection in US microcents for (second) east coast local setup testing point probing East Coast EC2 regions.

Overall, as expected, the per-connection cost is lowest intra-cloud (390-430 US microcents) and highest for connections from the outside (500-1000 US microcents). Inter-region costs are in between at around 400-490 US microcents.