

NATIONAL RECONNAISSANCE OFFICE

Historical AWS Pricing Trends

Andrew Kicinski, Hassan Souiri

Centauri

NRO/Cost and Acquisition Assessment Group (CAAG)

ICEAA May 2019




SUPRA ET ULTRA



Introduction

- ✦ Cost and Acquisition Assessment Group (CAAG)
 - ✦ Critical role: Agency Cost Positions – “how much will it cost?”
 - ✦ Key responsibility: provide independent cost and schedule analyses for program planning, budget, acquisition decisions and design reviews
- ✦ NRO utilizes cloud services provided by Amazon Web Services (AWS)
 - ✦ Dedicated region for Intelligence Community (IC)
 - ✦ Operating since late 2014
 - ✦ Utilization and therefore cost have been steadily increasing







Study Outline

- + Task: “To find and compile commercial AWS price history and see the pattern of price decreases in order to estimate future pricing”
 - + Storage – Simple Storage Service (S3)
 - + Compute – Elastic Compute Cloud (EC2)
- + AWS prides itself on offering periodical cuts to service pricing
 - + **66** total price reductions to commercial AWS since 2006
- + Multiple factors influence cloud pricing

- + Competition
- + Economies of scale
- + Hardware costs
- + License costs



Bureau of Labor Statistics (BLS) Producer Price Index (PPI) – Computer Storage Devices




How can 12 years of AWS pricing data be leveraged to model future prices?



Study Impact

- ✦ Just as omitting inflation impacts can lead to a large estimating error (under-estimate), ignoring cloud service price decreases will likely produce an over-estimate



Example with \$100 first year cost and constant usage

- ✦ Prior to the AWS Price Study, CAAG estimates assumed annual cloud price decreases based on BLS metrics and SME judgement

Study defends application of annual price reduction factors with rigorous analysis



Approach

- ✦ Collect historical commercial AWS S3 and EC2 pricing
- ✦ Normalize pricing for analysis
- ✦ Develop and evaluate candidate models to estimate future cloud service prices
- ✦ Calibrate/tailor model to AWS region utilized by NRO
 - ✦ Note: this presentation does not share the NRO model, but calibrates the commercial AWS results to AWS GovCloud using the same approach



Data Collection

- ✦ Collected raw data from Amazon to retrieve 2015-2018 pricing information
- ✦ Used archival website to find archived AWS S3 pricing dated back to 2006 and EC2 pricing to 2010
- ✦ Compiled the EC2 data by OS, Term Type and Product Family
- ✦ Normalized the data in order to find the annual percent price change
- ✦ Narrowed dataset to service options most relevant to NRO



S3 Model Assumptions

- ✦ Due to the tiered S3 pricing structure, one storage level had to be assumed to analyze costs over time
- ✦ Based on internal analysis, models assumed 7.11 PB

Months between price change		32	14	10	15	10	14	2	32	17
Effective Date		3/14/2006	11/1/2008	1/1/2010	11/1/2010	2/1/2012	12/1/2012	2/1/2014	4/1/2014	12/1/2016
Storage	0-1 TB	0.15	0.15	0.15	0.14	0.125	0.095	0.085	0.03	0.023
	1-50 TB	0.15	0.15	0.15	0.125	0.11	0.08	0.075	0.0295	0.023
	50-100 TB	0.15	0.14	0.14	0.11	0.095	0.07	0.06	0.029	0.022
	100-500 TB	0.15	0.13	0.13	0.11	0.095	0.07	0.06	0.029	0.022
	500-1000 TB	0.15	0.12	0.105	0.095	0.09	0.065	0.055	0.0285	0.021
	1000-5000 TB	0.15	0.12	0.08	0.08	0.08	0.06	0.051	0.028	0.021
	5000 TB or More	0.15	0.12	0.055	0.055	0.055	0.055	0.043	0.0275	0.021

Storage (TB)	Effective Price (\$ per GB/Month)									
Enter storage level -> 7,110	0.1500	0.1209	0.0781	0.0759	0.0744	0.0596	0.0496	0.0280	0.0211	





Calculate Annual Price Change

- ✦ Compare today's price to the S3 introductory price and find annual percent change from S3 introduction to today

$$\left(\frac{p_{today}}{p_0} \right)^{\left[\frac{1}{t_0 - t_{today}} \right]} - 1 = -14.9\%$$


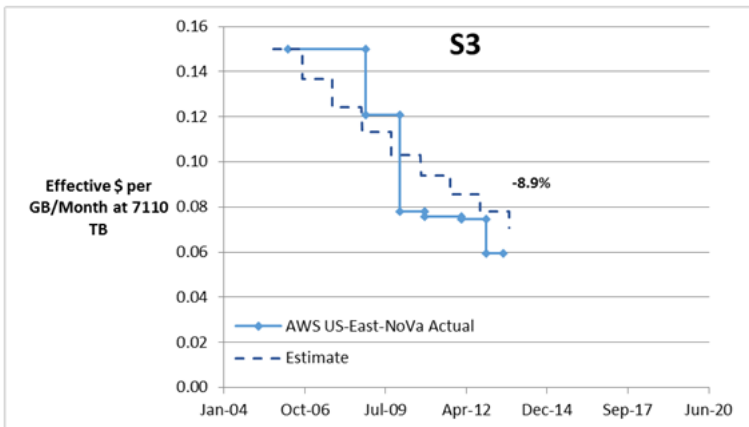

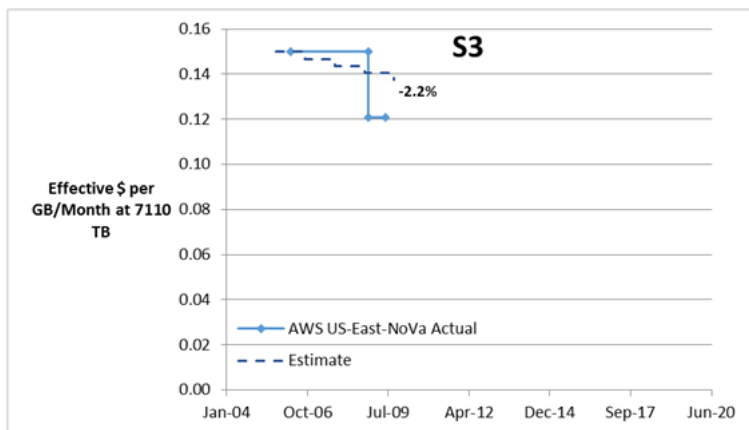
- ✦ Prices would be underestimated in almost the entire timeframe
- ✦ Next the underestimate bias was addressed by not constraining the introductory price
- ✦ This model resulted in an estimated 2006 price 25% higher than the actual price





Constrain Model Area and p_0

- ✦ Area between est. and act. curves constrained to 0; p_0 constrained to 0.15
- ✦ This method is always time biased (underestimates beginning years; overestimates later years)





Autoregressive (AR(p)) Time Series Model

- + To investigate an alternative method where the current price is effectively constrained to its actual value, an Autoregressive (AR) time series model was considered.
- + The AR model predicts a month's price based on a certain number of previous month's prices (referred to as order p , where p is the number of previous months considered).
- + The resulting model coefficients produce an exponential curve that predicts future prices.
- + An annualized factor is then calculated from the continuous exponential curve to produce a step function.

AR(p) Model

$$Z_t = c + \sum_{i=1}^p \varphi_i Z_{t-i} + \varepsilon_t$$

Z_t = price at time t

c = constant, set equal to zero for our models

p = order of model; number of previous months considered

φ = parameters or weights of the model

ε = error term



AR Model Order Selection

- + Time series models are typically fit to data that sporadically increase and decrease, such as stock market indices, home prices or the unemployment rate
- + Given the monotonic behavior of S3 prices and near-monotonic behavior of EC2 prices, this price analysis differs from other time series models
- + The table below shows the number of months between price reductions and the number of months between two consecutive price reductions


Months Between Price Reductions	32	14	10	15	10	14	2	32
Months Between Consecutive Price Reductions	46	24	25	25	24	16	34	

- + There was one price reduction that was followed two months later by another price reduction. Otherwise, all successive price reductions were separated by at least ten months.
- + AR(10), AR(14) and AR(15) seem like logical choices given the distances between price reductions.



Commercial AWS S3 AR(10) Model


- ✦ The time series model fits the actual cost function very closely, as expected due to the stable, non-increasing behavior of the cost function.
- ✦ Predicts 14.3% annual price reductions
- ✦ Reverse fitting the exponential curve appears to fit the general shape of the data, but generally underestimates actual cost





Commercial AWS S3 – Storage Prices

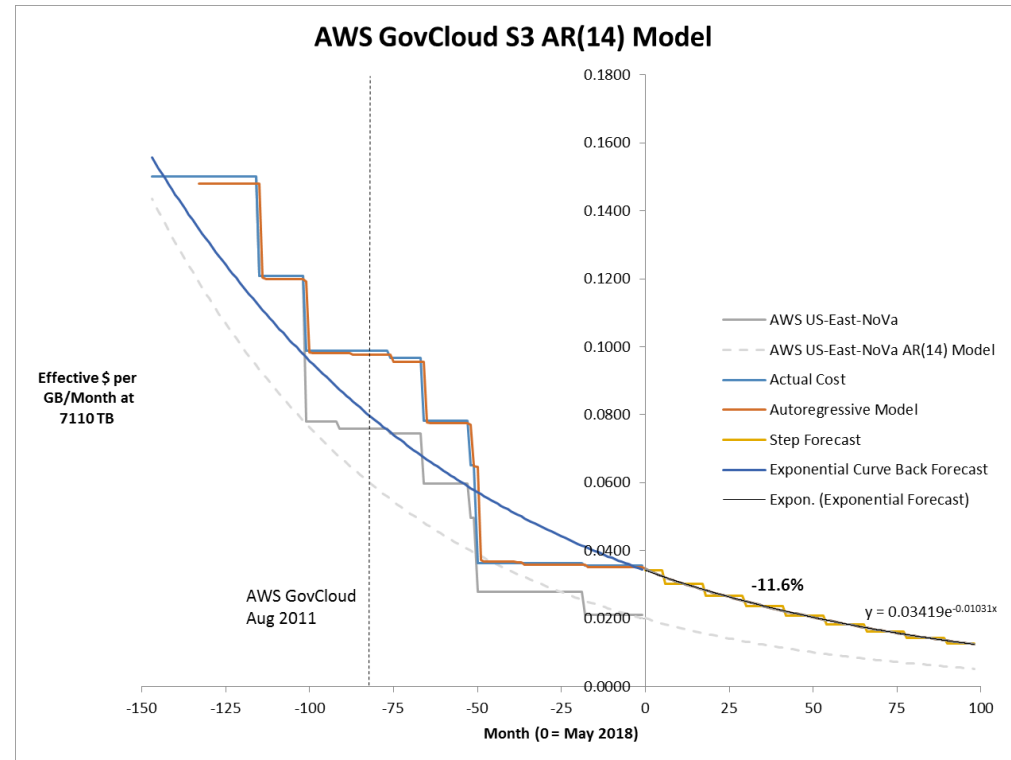
- + AR(14) model was considered using the same constraints and objectives as the AR(10) model
- + The back forecasted exponential curve is slightly higher than in the AR(10) model and represents a closer fit to the actual historical costs.
- + The model includes a longer price history covering more price reductions at some points and the back forecast curve aligning closer to the data, the AR(14) model is preferred over the AR(10) model.
- + 14 month autoregressive (AR) time series model predicts 14.9% annual price reductions on commercial S3





AWS GovCloud S3 Model

- + Commercial AWS faces more pressure from competition to reduce prices and benefits from larger economies of scale than AWS GovCloud
- + GovCloud S3 was introduced in August 2011
- + 14 month AR hybrid model generated using commercial AWS S3 pricing up to August '11 and GovCloud S3 pricing since
- + Model forecasts a future GovCloud S3 annual price decrease of 11.6%

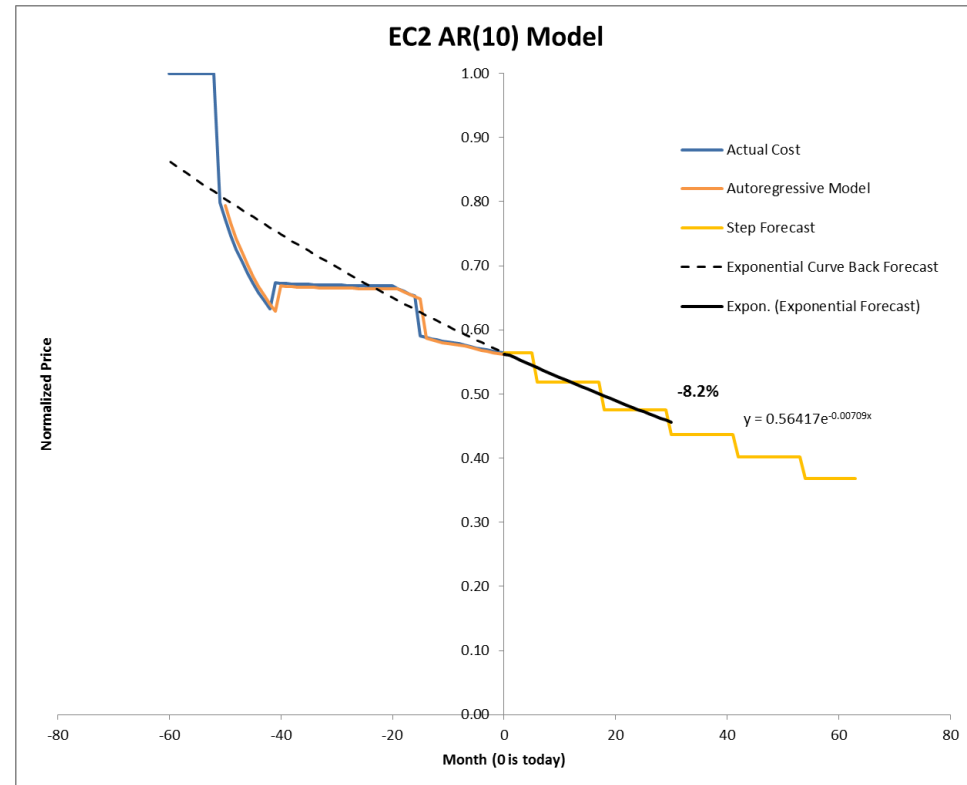


Model result of 11.6% is reasonable given BLS PPI for storage curve (11.4%)



Commercial AWS EC2 – Compute Prices

- + AWS EC2 offers thousands of combinations of service offerings
 - + e.g. m5.xlarge Linux 1 Year Standard Reserved All Upfront
 - + Analysis focused on popular services used
- + AWS EC2 prices have decreased 41% over the past five years (11% annually)
 - + Includes explicit service price reductions and new generation compute instances at lower prices compared to predecessors
- + 10 month AR proved best fit of AWS EC2 pricing data, forecasting 8.2% annual price reductions

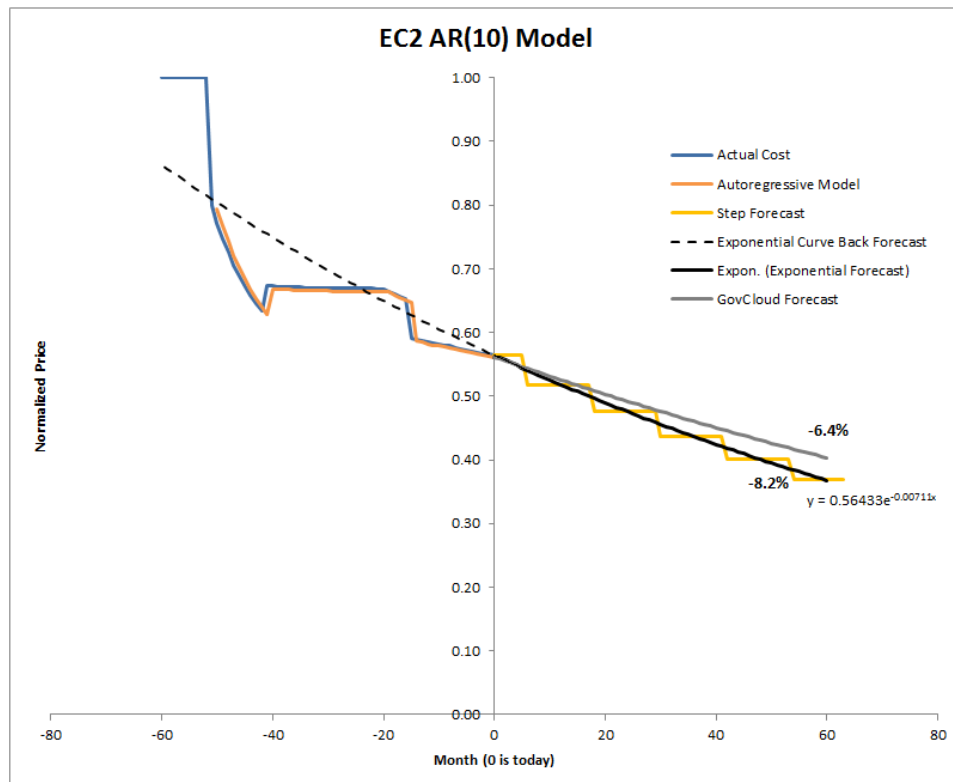




AWS GovCloud EC2 Model

- + To generate forecasted GovCloud EC2 prices, we apply the GovCloud to AWS S3 model ratio to the commercial EC2 model
 - + Results in projected annual price reductions of 6.4% on GovCloud EC2

$$+ \frac{11.6\%}{14.9\%} * 8.2\% = 6.4\%$$





Summary

- ✦ Commercial AWS has seen significant price reductions over the past decade across storage and compute services
- ✦ AWS GovCloud realizes similar effects as the commercial region that contribute to price decreases, namely economies of scale, improvements in hardware technology and need to be competitive in an active market
- ✦ Price reductions and new billing options that lower costs have been introduced in the more limited history of GovCloud
- ✦ Cost estimates and budget forecasts must model both price decreases and usage growth over time
- ✦ Prices will be monitored and models regularly adjusted

CAAG models forecast annual price decreases of 11.6% for S3 and 6.4% for EC2



Questions?



Thank you!

Hassan Souiri

souiriha@nro.mil

Hassan.Souiri@centauricorp.com

571-304-8902

NATIONAL RECONNAISSANCE OFFICE

SUPRA ET ULTRA

