



Tributary: spot-dancing for elastic services with latency SLOs

Aaron Harlap and Andrew Chung, *Carnegie Mellon University*; Alexey Tumanov, *UC Berkeley*; Gregory R. Ganger and Phillip B. Gibbons, *Carnegie Mellon University*

<https://www.usenix.org/conference/atc18/presentation/harlap>

**This paper is included in the Proceedings of the
2018 USENIX Annual Technical Conference (USENIX ATC '18).**

July 11–13, 2018 • Boston, MA, USA

ISBN 978-1-939133-02-1

**Open access to the Proceedings of the
2018 USENIX Annual Technical Conference
is sponsored by USENIX.**

Tributary: spot-dancing for elastic services with latency SLOs

Aaron Harlap^{§*} Andrew Chung^{§*} Alexey Tumanov[†]

Gregory R. Ganger[§] Phillip B. Gibbons[§]

[§]Carnegie Mellon University [†]UC Berkeley

Abstract

The Tributary elastic control system embraces the uncertain nature of transient cloud resources, such as AWS spot instances, to manage elastic services with latency SLOs more robustly and more cost-effectively. Such resources are available at lower cost, but with the proviso that they can be preempted *en masse*, making them risky to rely upon for business-critical services. Tributary creates models of preemption likelihood and exploits the partial independence among different resource offerings, selecting collections of resource allocations that satisfy SLO requirements and adjusting them over time, as client workloads change. Although Tributary’s collections are often larger than required in the absence of preemptions, they are cheaper because of both lower spot costs and partial refunds for preempted resources. At the same time, the often-larger sets allow unexpected workload bursts to be absorbed without SLO violation. Over a range of web service workloads, we find that Tributary reduces cost for achieving a given SLO by 81–86% compared to traditional scaling on non-preemptible resources, and by 47–62% compared to the high-risk approach of the same scaling with spot resources.

1 Introduction

Elastic web services have been a cloud computing staple from the beginning, adaptively scaling the number of machines used over time based on time-varying client workloads. Generally, an adaptive scaling policy seeks to use just the number of machines required to achieve its *Service Level Objectives (SLOs)*, which are commonly focused on response latency and ensuring that a given percentage (*e.g.*, 95%) of requests are responded to in under a given amount of time [17, 28, 19]. Too many machines results in unnecessary cost, and too few results in excess customer dissatisfaction. As such, much research and development has focused on doing this well [20, 14, 11, 12, 26].

Elastic service scaling schemes generally assume independent and infrequent failures, which is a relatively safe assumption for high-priority allocations in private clouds and *non-preemptible* allocations in public clouds (*e.g.*, on-demand instances in AWS EC2 [3]). This as-

sumption enables scaling schemes to focus on client workload and server responsiveness variations in determining changes to the number of machines needed to meet SLOs.

Modern clouds also offer transient, *preemptible* resources (*e.g.*, EC2 Spot Instances [1]) at a discount of 70–80% [6], creating an opportunity for cheaper service deployments. But, simply using standard scaling schemes fails to address the risks associated with such resources. Namely, preemptions should be expected to be more frequent than failures and, more importantly, preemptions often occur in bulk. Akin to co-occurring failures, bulk preemptions can cause traditional scaling schemes to have sizable gaps in SLO attainment.

This paper describes Tributary, a new elastic control system that exploits transient, preemptible resources to reduce cost and increase robustness to unexpected workload bursts. Tributary explicitly recognizes the bulk preemption risk, and it exploits the fact that preemptions are often not highly correlated across different pools of resources in heterogeneous clouds. For example, in AWS EC2, there is a separate spot market for each instance type in each availability zone, and researchers have noted that they often move independently: while preemptions within each spot market are correlated, across spot markets they are not [16]. To safely use preemptible resources, Tributary acquires collections of resources drawn from multiple pools, modified as resource prices change and preemptions occur, while endeavoring to ensure that no single bulk preemption would cause SLO violation. We refer to this dynamic use of multiple preemptible resource pools as *spot-dancing*.

AcquireMgr is Tributary’s component that decides the resource collection’s makeup. It works with any traditional scaling policy that determines (reactively or predictively) how many cores or machines are needed for each successive period of time, based on client load variation. AcquireMgr decides *which* instances will provide sufficient likelihood of meeting each time period’s target at the lowest expected cost. Its probabilistic algorithm combines resource cost and preemption probability predictions for each pool to decide how many resources to include from each pool, and at what price to bid for any new resources (relative to the current market price).

*Equal contribution

Given that a preemption occurs when a market's spot price exceeds the bid price given at resource acquisition time, AcquireMgr can affect the preemption probability via the delta between its bid price and the current price, informed by historical pricing trends. In our implementation, which is specialized to AWS EC2, the predictions use machine learning (ML) models trained on historical EC2 Spot Price data. The expected cost of the computation takes into account EC2's policy of partial refunds for preempted instances, which often results in AcquireMgr choosing high-risk instances and achieving even bigger savings than just the discount for preemptibility.

In addition to the expected cost savings, Tributary's spot-dancing provides a burst tolerance benefit. Any elastic control scheme has some reaction delay between an unexpected burst and any resulting addition of resources, which can cause SLO violations. Because Tributary's resource collection is almost always bigger than the scaling policy's most recent target in order to accommodate bulk preemptions, extra resources are often available to handle unexpected bursts. Of course, traditional elastic control schemes can also acquire extra resources as a buffer against bursts, but only at a cost, whereas the extra resources when using Tributary are a bonus side-effect of AcquireMgr's robust cost savings scheme.

Results for four real-world web request arrival traces and real AWS EC2 spot market data demonstrate Tributary's cost savings and SLO benefits. For each of three popular scaling policies (one reactive and two predictive), Tributary's exploitation of AWS spot instances reduces cost by 81–86% compared to traditional scaling with on-demand instances for achieving a given SLO (e.g., 95% of requests below 1 second). Compared to unsafely using traditional scaling with spot instances (*AWS AutoScale* [2]) instead of on-demand instances, Tributary reduces cost by 47–62% for achieving a given SLO. Compared to other recent systems' policies for exploiting spot instances to reduce cost [24, 16], Tributary provides higher SLO attainment at significantly lower cost.

This paper makes four primary contributions. First, it describes Tributary, the first resource acquisition system that takes advantage of preemptible cloud resources for elastic services with latency SLOs. Second, it introduces AcquireMgr algorithms for composing resource collections of preemptible resources cost-effectively, exploiting the partial refund model of EC2's spot markets. Third, it introduces a new preemption prediction approach that our experiments with EC2 spot market price traces show is significantly more accurate than previous preemption predictors. Fourth, we show that Tributary's approach yields significant cost savings and robustness benefits relative to other state-of-the-art approaches.

2 Background and Related Work

Elastic services dynamically acquire and release machine resources to adapt to time-varying client load. We distinguish two aspects of elastic control, the *scaling policy* and the *resource acquisition scheme*. The scaling policy determines, at any point in time, *how many* resources the service needs in order to satisfy a given SLO. The resource acquisition scheme determines *which* resources should be allocated and, in some cases, aspects of how (e.g., bid price or priority level). This section discusses AWS EC2 spot instances and resource acquisition strategies to put Tributary and its new approach to resource acquisition into context.

2.1 Preemptible resources in AWS EC2

In addition to non-preemptible, or *reliable* resources, most cloud infrastructures offer preemptible resources as a way to increase utilization in their datacenters. Preemptible resources are made available, on a best-effort basis, at decreased cost (in for-pay settings) and/or at lower priority (in private settings). This subsection describes preemptible resources in AWS EC2, both to provide a concrete example and because Tributary and most related work specialize to EC2 behavior.

EC2 offers “on-demand instances”, which are reliable VMs billed at a flat per-second rate. EC2 also offers the same VM types as “spot instances”, which are preemptible but are usually billed at prices significantly lower (70% - 80%) than the corresponding on-demand price. EC2 may preempt spot instances at any time, thus presenting users with a trade-off between reliability (on-demand) and cost savings (spot).

There are several properties of the AWS EC2 spot market behavior that affect customer cost savings and the likelihood of instance preemption. (1) Each instance type in each availability zone has a unique AWS-controlled spot market associated with it, and AWS's spot markets are not truly free markets [9]. (2) Price movements among spot markets are not always correlated, even for the same instance type in a given region [23]. (3) Customers specify a bid in order to acquire a spot instance. The bid is the maximum price a customer is willing to pay for an instance in a specific spot market; once a bid is accepted by AWS, it cannot be modified. (4) A customer is billed the spot market price (not the bid price) for as long as the spot market price for the instance does not exceed the bid price or until the customer releases it voluntarily. (5) As of Oct 2nd, 2017, AWS charges for the usage of an EC2 instance up to the second, with one exception: if the spot market price of an instance exceeds the bid price during its first hour, the customer is refunded fully for its usage. No refund is given if the spot instance is revoked in any subsequent hour. We define the period where preemption makes the instance free

as the *preemption window*.

When using EC2 spot instances, the bidding strategy plays an important role in both cost and preemption probability. Many bidding strategies for EC2 spot instances have been studied [9, 33, 30]. The most popular strategy by far is to bid the on-demand price to minimize the odds of preemption [23, 21], since AWS charges the market price rather than the bid price.

2.2 Cloud Resource Acquisition Schemes

Given a target resource count from a scaling policy, a resource acquisition scheme decides *which* resources to acquire based on attributes of resources (e.g., bid price or priority level). Many elastic control systems assume that all available resources are equivalent, such as would be true in a homogeneous cluster, which makes the acquisition scheme trivial. But, some others address resource selection and bidding strategy aspects of multiple available options. Tributary’s AcquireMgr employs novel resource acquisition algorithms, and we discuss related work here.

AWS AutoScale [2] is a service provided by AWS that maintains the resource footprint according to the target determined by a scaling policy. At initialization time, if using on-demand instances, the user specifies an instance type and availability zone. Whenever the scaling target changes, AutoScale acquires or releases instances to reach the new target. If using spot instances, the user can use a so-called “spot fleet” [4] consisting of multiple instance type and availability zone options. In this case, the user configures AutoScale to use one of two strategies. The *lowestPrice* strategy will always select cheapest current spot price of the specified options. The *diversified* strategy will use an equal number of instances from each option. Tributary bids aggressively and diversifies based on predicted preemption rates and observed inter-market correlation, resulting in both higher SLO attainment and lower cost than AutoScale.

Kingfisher [26] uses a cost-aware resource acquisition scheme based on using integer linear programming to determine a service’s resource footprint among a heterogeneous set of non-preemptible instances with fixed prices. Tributary also selects from among heterogeneous options, but addresses the additional challenges and opportunities introduced by embracing preemptible transient resources. Several works have explored ways of selecting and using spot instances. HotSpot [27] is a resource container that allows an application to suspend and automatically migrate to the most cost-efficient spot instance. While HotSpot works for single-instance applications, it is not suitable for elastic services since its migrations are not coordinated and it does not address bulk preemptions.

SpotCheck [25] proposes two methods of selecting spot markets to acquire instances in while always bid-

ding at a configurable multiple of the spot instance’s corresponding on-demand price. The first method is *greedy cheapest-first*, which picks the cheapest spot market. The second method is *stability-first*, which chooses the most price-stable market based on past market price movement. SpotCheck relies on VM migration and hot spares (on-demand or otherwise) to address revocations, which incurs additional cost, while Tributary uses a diverse pool of spot instances to mitigate revocation risk.

BOSS [32] hosts key-value stores on spot instances by exploiting price differences across pools in different data-centers and creating an online algorithm to dynamically size pools within a constant bound of optimality. Tributary also constructs its resource footprint from different pools, within and possibly across data-centers. Whereas BOSS assumes non-changing storage capacity requirements, Tributary dynamically scales its resource footprint to maintain the specified latency SLO while adapting to changes in client workload.

Wang et al. [31] explore strategies to decide whether, in the face of changing application behavior, it is better to reserve discounted resources over longer periods or lease resources at normal rates on a shorter term basis. Their solution combines on-demand and “reserved” (long term rental at discount price) instances, neither of which are ever preempted by Amazon.

ExoSphere [24] is a virtual cluster framework for spot instances. Its instance acquisition scheme is based on market portfolio theory, relying on a specified risk averseness parameter (α). ExoSphere formulates the *return* of a spot instance acquisition as the difference between the on-demand cost and the expected cost based on past spot market prices. It then tries to maximize the return of a set of instance allocations with respect to risk, considering market correlations and α , determining the fraction of desired resources to allocate in each spot market being considered. For a given virtual cluster size, ExoSphere will acquire the corresponding number of instances from each market at the on-demand price. Unsurprisingly, since it was created for a different usage model, ExoSphere’s scheme is not a great fit for elastic services with latency SLOs. We implement ExoSphere’s scheme and show in Section 5.6 that Tributary achieves lower cost, because it bids aggressively (resulting in more preemptions), and higher SLO attainment, because it explicitly predicts preemptions and selects resource sets based on sufficient tolerance of bulk preemptions.

Proteus [16] is an elastic ML system that combines on-demand resources with aggressive bidding of spot resources to complete batch ML training jobs faster and cheaper. Rather than bidding the on-demand price, it bids close to market price and aggressively selects spot markets and bid prices that it predicts will result in preemption, in hopes of getting many partial hours of free re-

sources. The few on-demand resources are used to maintain a copy of the dynamic state as spot instances come and go, and acquisitions are made and used to scale the parallel computation whenever they would reduce the average cost per unit work. Although Tributary uses some of the same mindset (aggressive use of preemptible resources), elastic services with latency SLOs are different than batch processing jobs; elastic services have a target resource quantity for each point in time, and having fewer usually leads to SLO violations, while having more often provides no benefit. Unsurprisingly, therefore, we find that Proteus’s scheme is not a great fit for such services. We implement Proteus’s acquisition scheme and show in Section 5.6 that Tributary achieves much higher SLO attainment, because it understands the resource target and explicitly uses diversity to mitigate bulk preemption effects. Tributary also uses a new and much more accurate preemption predictor.

3 Elastic Control in Tributary

AcquireMgr is Tributary’s resource acquisition component, and its approach differentiates Tributary from previous elastic control systems. It is coupled with a scaling policy, any of many popular options, which provides the time-varying resource quantity target based on client load. *AcquireMgr* uses ML models to predict the preemption probability of resources and exploits the relative independence of AWS spot markets to account for potential bulk preemptions by acquiring a diverse mix of preemptible resources collectively expected to satisfy the user-specified latency SLO. This section describes how *AcquireMgr* composes the resource mix while targeting minimal cost.

Resource Acquisition. *AcquireMgr* interacts with AWS to request and acquire resources. To do so, *AcquireMgr* builds sets of request vectors. Each request vector specifies the instance type, availability zone, bid price, and number of instances to acquire. We call this an *allocation request*. An *allocation* is defined as a set of instances of the same type acquired at the same time and price. *AcquireMgr*’s *total footprint*, denoted with the variable A , is a set of such allocations. Resource acquisition decisions are made under four conditions: (1) a periodic (one-minute) clock event fires, (2) an allocation reaches the end of its preemption window, (3) the scaling policy specifies a change in resource requirement, and/or (4) a preemption occurs. We term these conditions *decision points*.

AcquireMgr abstracts away the resource type which is being optimized for. For the workloads described in this paper, virtual CPUs (VCPUs) are the bottleneck resource; however, it is possible to optimize for memory, network bandwidth, or other resource types instead. A service using Tributary provides its resource scaling

characteristics to *AcquireMgr* in the form of a *utility function* $v()$. This *utility function* maps the number of resources to the percentage of requests expected to meet the target latency, given the load on the web service. The shape of a utility function is service-specific and depends on how the service scales, for the expected load, with respect to the number of resources. In the simplest case where the web service is embarrassingly parallel, the utility function is linear with respect to the number of resources offered until 100% of the requests are expected to be satisfied, at which point the function turns into a horizontal line. As a concrete example, if an embarrassingly parallel service specifies that 100 instances are required to handle 10000 requests per second without any of the requests missing the target latency, a linear utility function will assume that 50 instances will allow the system to meet the target latency on 50% of the requests. Tributary allows applications to customize the utility function so as to accommodate the resource requirements of applications with various scaling characteristics.

In addition to providing $v()$, the service also provides the application’s target SLO in terms of a percentage of requests required to meet the target latency. By exposing the target SLO as a customizable input, Tributary allows the application to control the Cost-SLO tradeoff. Upon receiving this information, *AcquireMgr* acquires enough resources to meet SLO in expectation while optimizing for expected cost. In deciding which resources to acquire, *AcquireMgr* uses the prediction models described in Sec. 3.1 to predict the probability that each allocation would be preempted. Using these predictions, *AcquireMgr* can compute the expected cost and the expected utility of a set of allocations (Sec. 3.2). *AcquireMgr* greedily acquires allocations until the expected utility is greater than or equal to the SLO percentage requirement (Sec. 3.3).

3.1 Prediction Models

When acquiring spot instances on AWS, there are three configurable parameters that affect preemption probability: instance type, availability zone and bid price. This section describes the models used by *AcquireMgr* to predict allocation preemption probabilities.

Previous work [16] proposed taking the historical median probability of preemption based on the instance type, availability zone and bid price. This approach does not consider time of day, day of week, price fluctuations and several other factors that affect preemption probabilities. *AcquireMgr* trains ML models considering such features to predict resource reliability.

Training Data and Feature Engineering. The prediction models are trained ahead of time with data derived from AWS spot market price histories. Each sample in the training dataset is a *hypothetical bid*, and the

target variable, `preempted`, of our model is whether or not an instance acquired with the hypothetical bid is preempted before the end of its preemption window (1 hr). We use the following method to generate our data set: For each instance and *bid delta* (bid price above the market price with range $[0.00001, 0.2]$) we generate a set of hypothetical bids with the bid starting at a random point in the spot market history. For each bid, we look forward in the spot market price history. If the market price of the instance rises above the bid price at any point within the hour, we mark the sample as `preempted`. For each historical bid, we also record the ten prices immediately prior to the random starting point and their time-stamps.

To increase prediction accuracy, `AcquireMgr` engineers features from AWS spot market price histories. Our engineered features include: (1) Spot market price; (2) Average spot market price; (3) Bid delta; (4) Frequency of spot market price changing within past hour; (5) Magnitude of spot market price fluctuations within past two, ten, and thirty minutes; (6) Day of the week; (7) Time of day; (8) Whether the time of day falls within working hours (separate feature for all three time zones). These features allow `AcquireMgr` to construct a more complex prediction model, leading to a higher prediction accuracy (Sec. 5.7).

Model Design. To capture the temporal nature of the EC2 spot market, `AcquireMgr` uses a *Long Short-Term Memory Recurrent Neural Network (LSTM RNN)* to predict instance preemptions. The LSTM RNN is a popular model for workloads where the ordering of training examples is important to prediction accuracy [29]. Examples of such workloads include language modeling, machine translation, and stock market prediction. Unlike feed forward neural networks, LSTM models take previous inputs into account when classifying input data. Modeling the EC2 spot market as a *sequence* of events significantly improves prediction accuracy (Sec. 5.7). The output of the model is the probability of the resource being preempted within the hour.

3.2 AcquireMgr

To make decisions about which resources to acquire or release, `AcquireMgr` computes the expected cost and expected utility of the set of instances it is considering at each decision point. Calculations of the expected values are based on probabilities of preemption computed by `AcquireMgr`'s trained LSTM model. This section describes how `AcquireMgr` computes these values.

Definitions. To aid in discussion, we first define the notion of a *resource pool*. Each instance type in each availability zone forms its own *resource pool*—in the context of the EC2 spot instances, each such resource pool has its own spot market. Given a *set of allocations* A , where A is formulated as a jagged array, let A_i be defined as the i^{th} entry of A corresponding to an array of

A	Set of allocations as jagged array
A_i	Sorted array of allocations from resource pool i
$a_{i,j}$	Set of instances allocated from resource pool i
$\beta_{i,j}$	Probability that allocation $a_{i,j}$ is preempted
$t_{i,j}$	Time left in the preemption window for $a_{i,j}$
$k_{i,j}$	Number of instances in allocation $a_{i,j}$
$P_{i,j}$	Market price of allocation $a_{i,j}$
$p_{i,j}$	Bid price of allocation $a_{i,j}$
$size(y)$	Size of the major dimension of array y
$resc(y)$	Counts the total number of resources in y
λ_i	Regularization term for diversity
$P(R=r)$	Probability that r resources remain in A
$v(r)$	The utility of having r resources remain in A
V_A	The expected utility of a set of allocations A
C_A	Expected cost of a set of allocations (\$)

Table 1: Summary of parameters used by `AcquireMgr`

allocations from resource pool i sorted by bid price in ascending order. We define allocation $a_{i,j}$ as an allocation from resource pool i (i.e., $a_{i,j} \in A_i$) with the j^{th} lowest bid in that resource pool. We further denote $p_{i,j}$ as the bid price of allocation $a_{i,j}$, $\beta_{i,j}$ as the probability of preemption of allocation $a_{i,j}$, and $t_{i,j}$ as the time remaining in the preemption window for allocation $a_{i,j}$. Note that $p_{i,j} \geq p_{i,j-1}$, which also implies $\beta_{i,j-1} \geq \beta_{i,j}$. Finally, we define a *size*(A) function that returns the size of A 's major dimension. See Table 1 for symbol reference.

Expected Cost. The total expected cost for a given footprint A is calculated as the sum over the expected cost of individual allocations $C_A[a_{i,j}]$:

$$C_A = \sum_{i=1}^{size(A)} \sum_{j=1}^{size(A_i)} C_A[a_{i,j}] \quad (1)$$

`AcquireMgr` calculates the *expected* cost of an allocation by considering the probability of preemption within the preemption window $\beta_{i,j}$ for a given allocation $a_{i,j}$ at a given *bid delta*. There are exactly two possibilities: an allocation will either be preempted with probability $\beta_{i,j}$ or it will reach the end of its preemption window in the remaining $t_{i,j}$ minutes with probability $1 - \beta_{i,j}$, in which case we would voluntarily release the allocation. The expected cost can then be written down as:

$$C_A[a_{i,j}] = (1 - \beta_{i,j}) * P_{i,j} * k_{i,j} * t_{i,j} + \beta_{i,j} * 0 * k_{i,j} * t_{i,j} \quad (2)$$

where $k_{i,j}$ is the number of instances in the allocation, and $P_{i,j}$ is the market price for instance of type i at the time of acquisition.

Expected Utility. In addition to computing expected cost for a set of allocations, `AcquireMgr` computes the *expected utility* for a set of allocations. The *expected utility* is the expected percentage of requests that will meet the latency target given the set of allocations A . Expected utility takes into account the probability of allocation preemptions, providing `AcquireMgr` with a metric

for quantifying the expected contribution that each allocation should make to meet the resource target. The expected utility V_A of the set of allocations A is calculated as follows:

$$V_A = \sum_{r=0}^{resc(A)} P(R=r) * v(r) \quad (3)$$

where $P(R)$ is the probability mass function of the discrete random variable R that denotes the number of resources not preempted within the next hour, v is the utility function provided by the service, and $resc(A)$ is the function that reports the number of resources in a set of allocations A . $resc(A)$ computes the *total amount of resources* in A , while $size(A)$ only computes the *size* of A 's major dimension.

Eq. 3 computes the expected utility over the next hour given a workload, as though Tributary just bid for all its allocations. This works, even though there will usually be complex overlapping expiration windows across an hour, because it only needs to hold true until recomputed at the next decision point, which is never more than a minute away. To derive $P(R)$, AcquireMgr starts off with the original set of allocations A and generates all possible subsets of A . Each possible subset $S \subseteq A$, S marks some allocations in A as preempted ($\in S$) and the remaining allocations as not preempted ($\notin S$). To formalize the notion, we define the indicator variable $d_{i,j}$, where $d_{i,j} = 1$ if allocation $a_{i,j} \in S$ and $d_{i,j} = 0$ otherwise.

To compute the probability of S being the set of preempted resources ($P(S)$), AcquireMgr separates all allocations by resource pools, as each resource pool within AWS has its own spot market. We leverage the following localizing property. Within each resource pool A_i , the probability of preempting an allocation $a_{i,j}$ is only dependent on whether the allocation with the next lowest bid price, $a_{i,j-1}$, in the same resource pool is preempted. Note that $P(a_{i,1}) = \beta_{i,1}$ for allocation $a_{i,1}$ for all resource pools i . Consider two allocations $a_{i,j}, a_{i,j-1} \in A$ from resource pool A_i . We observe the following properties: (1) $a_{i,j}$ cannot be preempted unless $a_{i,j-1}$ is preempted, (2) the probability that both $a_{i,j}$ and $a_{i,j-1}$ are preempted is the probability that $a_{i,j}$ is preempted, and (3) the probability that $a_{i,j}$ is preempted without $a_{i,j-1}$ being preempted is 0. With Bayes' Rule, we observe that:

$$P(a_{i,j}|a_{i,j-1}) = \frac{P(a_{i,j} \wedge a_{i,j-1})}{P(a_{i,j-1})} = \frac{\beta_{i,j}}{\beta_{i,j-1}}. \quad (4)$$

Thus, for an allocation $a_{i,j}$ given subset $S \subseteq A$,

$$P(a_{i,j}|a_{i,j-1}) = \begin{cases} 0 & \text{if allocation } a_{i,j-1} \notin S, \\ \beta_{i,j}/\beta_{i,j-1} & \text{else.} \end{cases} \quad (5)$$

Tributary further introduces a regularization term λ_i to encourage bidding in markets with low correlation. Having instances spread across lowly correlated markets

is important for avoiding high-risk footprints. If the resource footprint has too many instances from correlated resource pools, Tributary becomes exposed to having too many resources being lost to a correlated price spike, potentially causing an SLO violation. In order obtain price correlation across spot markets, we periodically keep track of fix-sized moving windows of spot markets and compute the Pearson correlation between each pair of spot markets. When computing *expected utility*, Tributary increases an allocation in A_i 's probability of preemption $\beta_{i,j}$ by λ_i :

$$\lambda_i = \gamma * \sum_{l=1}^{size(A)} \rho_{i,l} * \frac{resc(A_i) + resc(A_l)}{2 * resc(A)} \quad (6)$$

where $\rho_{i,l}$ is the Pearson correlation between resource pools i and l , and $\gamma \in \mathbb{R} \geq 0$ is the configurable penalty multiplier. Essentially, we add a weighted penalty to an allocation based on its Pearson correlation scores with the rest of our resources in different resource pools. In our experiments, we set $\gamma = 0.01$. The regularization term leads to Tributary creating a diversified resource pool, thus reducing the probability that a significant portion of the resources are preempted simultaneously. Having a high probability of maintaining the majority of the resource pool at any point time, allows Tributary to avoid SLO violations with a high probability.

Let's denote $P(S)$ as the probability of S being the set of resources preempted from A . AcquireMgr computes it by taking the product of the conditional probability of each allocation having the outcome specified in S . If the allocation is preempted ($d_{i,j} = 1$) the conditional probability of the allocation being preempted ($P(a_{i,j}|a_{i,j-1})$) is used, otherwise ($d_{i,j} = 0$) the product uses the conditional probability of the allocation not being preempted ($1 - P(a_{i,j}|a_{i,j-1})$).


$$P(S) = \prod_{i=1}^{size(A)} \prod_{j=1}^{size(A_i)} \left(d_{i,j} * P(a_{i,j}|a_{i,j-1}) + (1 - d_{i,j}) * (1 - P(a_{i,j}|a_{i,j-1})) \right) \quad (7)$$

Finally, AcquireMgr formulates the probability of r resources remaining after preemption $P(R=r)$ (Eq. 3) as the sum of the probabilities of all sets S where the number of resources not preempted in S equals to r :

$$P(R=r) = \sum_{S \subseteq A, resc(S)=resc(A)-r} P(S) \quad (8)$$

which it uses to calculate the expected utility of a set of allocations A (Eq. 3).

Computational tractability. AcquireMgr's algorithm is exponentially computationally expensive as the number of spot markets considered increases. When considering more markets, it is possible to reduce computational complexity by grouping similar, correlated spot



(a) Legend (b) Tributary (c) AutoScale

Figure 1: Figures (b) and (c) show how Tributary and AutoScale handle a sample workload respectively. Figure (a) is the legend for (b) and (c), color-coding each allocation. The black dotted lines in (b) and (c) signify the request rates over time. At **minute 15**, the request rate unexpectedly spikes and AutoScale experiences “slow” requests until completing integration of additional resources with 3. Tributary, meanwhile, had extra resources meant to address preemption risk in C, providing a natural buffer of resources that is able to avoid “slow” requests during the spike. At **minute 35**, when the request rate decreases, Tributary terminates B, since it believes that B has the lowest probability of getting the free partial hour. It does not terminate D since it has a high probability of eviction and is likely to be free; it also does not terminate C since it needs to maintain resources. AutoScale, on the other hand, terminates both 2 and 3, incurring partial cost. At **minute 52**, the request rate increases and Tributary again benefits from the extra buffer while AutoScale misses its latency SLO. In this example, Tributary has less “slow” requests and achieves lower cost than AutoScale because AutoScale pays for 3 and for the partial hour for both 1 and 2 while Tributary only pays for A and the partial hour for B since C and D were preempted and incur no cost.

markets, and performing revocation analysis with a representative market. Although this would potentially decrease the precision of the preemption analysis, it would allow AcquireMgr to further improve performance by considering a larger number of markets.

3.3 Scaling Out

Resource Acquisition. When Tributary starts, the user specifies a *target SLO* in terms of percentage of requests that respond within a certain latency for Tributary to target. AcquireMgr uses this target SLO to acquire resources. At each decision point, AcquireMgr’s objective is to acquire resources until the expected utility θ_A is greater than or equal to the target SLO. If the expected utility is greater than or equal to the target SLO, no action is taken; otherwise, AcquireMgr computes the expected cost (Eq. 2) and utility of the current set of allocations (Eq. 3). AcquireMgr then calculates the missing number of resources (M) required to meet the target SLO and builds a set of possible allocations (Λ) that consists of allocations from different resource pools at different bid prices (from \$0.0001 to \$0.2 above the current price). For each possible allocation Λ_i , AcquireMgr records the new expected utility divided by the new expected cost of $A \cup \Lambda_i$, choosing the allocation Λ_{chosen} that maximizes this value. AcquireMgr continues to add possible allocations until it achieves the target SLO in expectation.

Buffers of Transient Resources. To accommodate potential resource preemptions, Tributary inherently acquires more than the required amount of resources if any of its allocations have a preemption probability greater than zero, which is always the case with spot instances. The amount of additional resources acquired depends on the target SLO and the probabilities of allocation pre-

emptions (Eq. 3). While the primary goal of these additional resources is to account for preemptions, they often have the added benefit handling unexpected increases in load. Experiments with Tributary show that these resource buffers both increase the fraction of requests meeting latency targets and decrease cost (Sec. 5.3).

3.4 Scaling In

Aside from preemptions, Tributary also tries to scale in voluntarily. As described earlier, each allocation is considered only for the duration of the preemption window. When an allocation reaches the end of its preemption window, it is terminated and replaced with a new allocation if required. When resource requirements decrease, Tributary considers terminating allocations for allocations least likely to be preempted. During this process Tributary chooses the allocation with the least time remaining in the hour, computes the expected utility θ_A without this allocation, and if it is greater than the target SLO, Tributary terminates the allocation. Tributary continues to try and terminate allocations as long as θ_A remains greater than the target SLO.

3.5 Example and Future Consideration

Example. Fig. 1 shows how Tributary and AutoScale handle a sample workload, including how the extra resources Tributary acquires to handle preemption events can also handle an unexpected request rate increase and how aggressive allocation selection can get some resources for free due to preemptions.

Future. Tributary lowers cost and meets SLO requirements by taking advantage of low-cost spot instances and uncorrelated prices across different spot instance markets. Mass adoption of systems like Tributary could

change these characteristics. While a detailed analysis of mass adoption’s potential effects on EC2 spot-markets is outside the scope of this paper, we evaluate the effects of two potential changes to the spot-market policies in Section 5.5.

4 Tributary Implementation

Figure 2 shows Tributary’s high-level system architecture. This section describes the main components, how they fit together, and how they interact with AWS.

Preemption Prediction Models. The *prediction models* are trained offline using TensorFlow [8] and deployed using Tensorflow Serving [7]. A separate model is used for each resource pool. To service run time predictions Tributary launches a *Prediction Serving Proxy* that receives all prediction queries from AcquireMgr, forwards them to their respective models, aggregates the results, and returns the predictions to AcquireMgr.

Resource Footprint Management. In Tributary, AcquireMgr takes primary responsibility for managing the resource footprint. AcquireMgr acquires instances, terminates instances, and monitors AWS for instance preemption notifications. AcquireMgr considers modifying the resource footprint at every *decision point*, and it follows the procedure described in Sec. 3.3 when additional resources are needed. Once AcquireMgr selects a set of instances to acquire, it sends instance requests to AWS via *boto.ec2* API calls. AWS responds with a set of spot request ids, which corresponds to the EC2 instances allocated to AcquireMgr. Once the instances are in a running state, AcquireMgr sends the instance ids associated with the new instances to Resource Manager. Instance removal follows a similar procedure.

Scaling Policy. The *Scaling Policy* component determines dynamic sizing of the resource target. Through a simple event-driven API, users can implement their own scaling policies that access metrics provided by the Monitoring Manager and specify the resource target.

Monitoring Manager (MonMgr). The *Monitoring Manager* orchestrates monitoring of service system resources. The Scaling Policy can register for metrics such as total number of requests and average CPU utilization of instances. The MonMgr queries requested metrics using AWS CloudWatch each *monitoring period* and forwards them to the scaling policy.

Resource Manager (ResMgr). The *Resource Manager* is a proxy for AcquireMgr. Using resource targets provided by the Scaling Policy, the ResMgr generates the utility function used by AcquireMgr to make resource acquisition decisions.¹ The ResMgr also receives instance allocations and termination notices from AcquireMgr and forwards them to the Service Manager.

¹Process of constructing the utility function is described in Sec. 5.2.




Figure 2: Tributary architecture.

5 Evaluation

This section evaluates Tributary’s effectiveness. The results support a number of important findings: (1) Tributary’s exploitation of AWS spot market instances reduces cost by 81%–86% compared to on-demand instances and simultaneously decrease SLO latency misses; (2) Compared to standard bidding policies for spot instances, Tributary reduces cost by up to 41% and decreases SLO latency misses by 31%–65%; (3) Compared to extending those standard policies to use enough extra (buffer) resources to match Tributary’s number of SLO latency misses, Tributary reduces cost by 47%–62%; (4) Tributary outperforms state-of-the-art resource managers in running elastic services; (5) Tributary’s preemption prediction models improve accuracy significantly, resulting in 37% lower cost than previous prediction approaches.

5.1 Experimental Setup

Experimental Platform. We report results for use of three AWS EC2 spot instance types: *c4.large*, *c4.xlarge*, and *c4.2xlarge*. The results correspond to the *us-west-2* region, which consists of three availability zones. Using the three instance types in each availability zone, our experiments involve nine resource pools.

Workload. The simulated workload uses a real-world trace for request arrival times, with each request consisting of the derivation of the PBKDF2 [18] key of a password. The calculation of a PBKDF2 key is CPU-heavy, with no network overhead and minimal memory overhead. With the CPU performance being the bottleneck, the resource requirement can be characterized in requests-per-second-per-VCPU.

Environment. In the simulation framework, each instance is characterized with a number of VCPUs, and the request processing time is configured to the measured time for one request on an EC2 instance (≈ 100 ms). Each instance server maintains a queue of requests, and we simulate the queueing effects using the discrete event simulation library SimPy [22]. The simulation framework takes into account resource start-up time, with newly acquired instances not able to service requests for two hundred seconds following their launch.

SLO and Scaling. The target service latency is set to




Figure 3: Traces used in system evaluation.

one second, and we verified on EC2 that a VCPU can handle roughly 10 requests per second without violating the latency target. So, the requests-per-second-per-VCPU is ten, and the queue size per server instance is ten times the number of VCPUs in the instance. Tributary is not overly sensitive to the target latency setting.

Traces. We use four real-world request arrival traces with differing characteristics. *Berkeley* is from the Berkeley Home IP proxy service and *ClarkNet* is from the ClarkNet ISP’s HTTP servers [10]. Both exhibit a periodic, diurnal pattern. We use the first 2000 minutes of these two traces, which covers an entire period. *WITS* is a sampled trace from the Waikato Internet Traffic Storage (WITS) [15]. The trace lasts for roughly a day, from April 6th to April 7th of the year 2000. This trace exhibits large variation of request rates throughout the day, as can be seen in Fig. 3b. *WorldCup98* is the arrival trace of the workload on the 1998 FIFA World Cup HTTP Servers [10] on day 75 of the World Cup. All traces are scaled to have an average of 125 requests per second in order to generate sufficient load for the experiments.

5.2 Scaling Policies Evaluated

We implement three popular scaling policies: *Reactive*, *Predictive Moving Window Average (MWA)*, and *Predictive Linear Regression (LR)* to evaluate our system. The utility function provided by the service is linear for all three policies. We make this assumption since our workload characteristic is embarrassingly parallel — if a workload exhibits different scaling characteristics, a different utility function can be employed.

The *Reactive Policy* scales out immediately when demand reported by the MonMgr is greater than what the available resources are able to handle. It scales in slowly (only after three minutes of low demand), as recommended by Gandhi et al. [12], to prevent premature scale-in in case the demand fluctuates widely in a short period of time. The *MWA Policy* maintains a sliding window of a fixed size, with each window entry consisting of the number of requests received in each monitoring period. The policy takes the average of the window entries to predict the number of requests on the next monitoring period. The policy then adjusts the utility and scaling functions according to the predicted number of requests, and reports the updated functions to the ResMgr to scale in expectation of future requests. The *LR Policy* also maintains a sliding window of a fixed size, but rather than us-

ing the average in the window for prediction, the policy performs linear regression on data points in the window to estimate the expected number of requests in the next monitoring period. Our experiments show that regardless of the scaling policy used, Tributary beats its competitors in both meeting the service latency target and cost.

5.3 Improvements with Tributary

Here, we evaluate Tributary’s ability to reduce cost and latency target misses against AutoScale.

AWS Autoscale. AWS AutoScale (Sec. 2.2) as offered by Amazon only supports the simplest reactive scaling policies. To provide better comparison between approaches, we implement the AWS AutoScale resource acquisition algorithm as closely as possible according to its documentation [2] and integrate it with Tributary’s SvcMgr to work with its more powerful scaling policies. From here on, mentions of *AutoScale* refer to our implementation of AWS AutoScale. AutoScale is the equivalent of the AcquireMgr component of Tributary. The default AutoScale algorithm with spot instances bids for the lowest market-priced spot instance at the on-demand price upon resource requests by the scaling policy. In addition, AutoScale terminates resources as soon as the resource requirements are lowered, choosing to terminate resources that are most expensive at the moment.

Methodology and Terminology. To achieve fair comparisons across a wide range of data points, we perform cost analysis with simulations using historical spot market traces. Using traces allows us to test different approaches on the same period of market data and to get a better picture of the expected behavior of the system in a shorter amount of time. For each request arrival trace (Sec. 5.1) and resource acquisition approach, we present the average cost and percentage of “slow” requests over trace requests across ten randomly chosen day/time starting points between January 23, 2017 and March 23, 2017 in the *us-west-2* region. From here on, we define a “slow” request as a request that does not meet the latency target and the percentage of “slow” requests as the percentage of “slow” requests over all requests in a single trace.²

Cost Savings and Service Latency Improvements. Fig. 4 shows the cost savings and percentage of “slow” requests for the *ClarkNet* trace. The cost savings are normalized against running Tributary on on-demand resources. The results demonstrate that Tributary reduces cost and “slow” requests for all three scaling policies. Cost savings are $\approx 85\%$ compared to on-demand resources. For the *ClarkNet* trace, Tributary reduces cost by 36%, 24% and 21% compared to to AutoScale for the *Reactive*, *Predictive-LR* and *Predictive-MWA* scaling policies, respectively. Compared to AutoScale, Tributary

²Prediction models were trained on data from 06/06/16 – 01/22/17.




Figure 4: Cost savings (red) and percentage of “slow” requests (blue) for the *ClarkNet* trace.

reduces “slow” requests by 72%, 61% and 64%, respectively, for the three scaling policies.

In order to decrease the number “slow” requests, popular scaling policies are often configured to provision more resources than immediately necessary to handle unexpected increases in load. It is common to specify the resource buffer as a percentage of the expected resource requirement. For example, with a buffer of 50%, 15 resources (e.g., VCPUs) would be acquired rather than the projected 10. AutoScale+Buffer shows the cost of provisioning AutoScale with a large enough buffer such that its number of “slow” requests matches that of Tributary. Tributary reduces cost by 61%, 56% and 57% compared to AutoScale+Buffer for the three scaling policies.

The cost savings for Tributary on the *Berkeley* trace relative to AutoScale are similar to those on the *ClarkNet* trace, but the reduction in percentage of “slow” requests increases. This difference in performance is due to differing characteristics of the two traces — the *ClarkNet* trace experiences more minute-to-minute volatility in request rate compared to the *Berkeley* trace. We observe similar levels of cost reductions and reduction in “slow” requests on the *WITS* and *WorldCup98* traces, results for *WITS* are shown in Tables 2. Compared to AutoScale+Buffer, Tributary decreased costs by 47–62% across all traces.

Scaling Policy	Cost Saving	“Slow” request Reduction
Reactive	37%	31%
Predictive-LR	33%	50%
Predictive-MWA	29%	51%

Table 2: Cost and “slow” request improvements for Tributary compared to AutoScale for the *WITS* trace

Attribution of Benefits. Tributary’s superior performance arises from several factors. Much of the reduction in cost compared to AutoScale is due to Tributary’s ability to get free instance hours. *Free instance hours* occur when an allocation does useful work but is preempted by AWS before the end of a preemption window. The user receives a refund for the partial hour, which means that any work done by the allocation in the preemption window comes at no cost to the user. Tributary takes the probability of getting free instance hours into account when computing the expected cost of allocations (Eq. 1), often acquiring resources that provide higher opportunities for free instance hours.

Another factor in Tributary’s lower cost is its ability to remove allocations that are not likely to be pre-

empted when demand drops. When resource demand decreases, Tributary terminates instances that are least likely to be preempted, thus lowering the expected cost of its resource footprint. The reductions in “slow” requests arise from the buffer of resources acquired by Tributary (Sec. 3.3). When acquiring instances, AcquireMgr estimates their probability of preemption. Unless all allocations have a preemption probability of zero, which never occurs for spot instances, Tributary acquires more resources than specified by the scaling policy. The primary goal of the additional resources is to ensure that, when Tributary experiences preemption events, it still has at least the specified number of resources in expectation. The additional resources also provide a secondary benefit by handling some or all of unexpected bursts of requests that exceed the load expected by the scaling policy. The cost of these additional resources is commonly offset by free instance hours; indeed, the extra resources are acquired to cope with preemptions.

5.4 Risk Mitigation

A key feature of Tributary is that it encourages instance diversification, *i.e.*, acquiring instances from mostly independent resource pools (Sec. 3.2). The default AutoScale policy is the lowest-price policy, which does not take diversification into account when acquiring instances; instead, it acquires the cheapest instance. Illustrated in Fig. 1, Tributary acquires different types of instances in different availability zones, while AutoScale acquires instances of the same type (all red). Diversifying across resource pools is important, because each has an independent spot market, avoiding highly correlated allocation preemptions within a single instance market. Acquiring too much from a single pool, as often occurs with AutoScale, creates a high risk of SLO violation when preemption events occur (e.g., if the red allocation in Fig. 1c was preempted prior to minute 35).

In our experiments, we found it to be very rare for market prices to rise above on-demand prices, meaning that AutoScale rarely experiences preemption events. However, when examining past EC2 spot market traces and other availability zones, we found it to be significantly more common for the market price to rise above the on-demand price, thus preempting AutoScale instances.³

³ From 01/23/17–03/20/17, the market price rose above the on-demand price 0 times for the *c4.2xlarge* instance type in *us-west-2*. From 11/1/16–01/22/17, it happened 1073 times.




Figure 5: Comparing to ExoSphere and Proteus. Predictive-MWA results not shown but similar.

Since Amazon charges users the market price and not the bid price, it is possible that Amazon may once again preempt instances bidding the on-demand price with regularity—a phenomenon we recently observed in the *us-east* availability zones. Thus, AutoScale’s resource acquisition approach is riskier for services with latency SLOs on spot machines.

Cost of Diversified AutoScale. In addition to the default AutoScale policy which acquires the lowest-priced instance, AWS also offers a diversified AutoScale policy that starts instances from a diverse set of resource pools [4]. Acquiring instances from different spot markets reduces preemption risks, but our experiments showed that it increases cost by 8%–12% compared to the lowest-price AutoScale policy. Compared to Tributary, which diversifies across spot markets intelligently, we found that a diversified AutoScale policy cost 68% more to achieve the same number of “slow” requests for the *reactive* scaling policy on the *ClarkNet* trace.

5.5 Pricing Model Discussion

Our experimental results are based on current AWS EC2 billing policies, as described Section 2.1. This section discusses how Tributary would function under two potential changes to the billing model: (1) elimination of preemption refunds, (2) institution of a free market.

Elimination of preemption refunds. If Amazon eliminates refunds when the market price exceeds bid price during the first hours of usage, Tributary would lose incentive to bid close to market price. Tributary’s model would capture this change by setting β in Eq. 2 to zero. With higher bids, Tributary would acquire fewer resources because preemption would be less likely. The amount of resources acquired would still exceed the amount of resources required as they would still have non-zero preemption probabilities.

Although Tributary extracts significant benefit from the refunds, it still outperforms AutoScale without it. For example, in a simulation with this billing model modification, Tributary still reduces cost by 31% compared to AutoScale with sufficient buffer to match numbers of “slow” requests, for the *Clarknet* trace using the *reactive* scaling policy. As expected, Tributary continues to meet SLOs with high likelihood, as it continues to diversify its resource pool and acquire buffers of resources (albeit smaller ones) to account for preemption events.

Free market behavior. In its current design, the AWS EC2 spot markets do not behave as free markets [9]. Cus-

tomers specify their bid prices for a given resource, but generally do not pay that amount. Instead, a customer is billed according to the EC2-determined spot price for that resource. It is possible, perhaps even likely as the spot market becomes widely popular, that AWS will transition toward a billing policy in which users are charged their *bid price*, instead of the *market price*, and prices move based on supply and demand rather than unknown seller policies. This change would render the commonly used strategy of bidding far above the market price (e.g., bidding the on-demand price) obsolete. Tributary’s behavior would not change significantly, as it already often sets bid prices close to market prices and explicitly considers revocation risks, and we believe it would therefore outperform other approaches by even larger margins.

5.6 Comparing to State of the Art

This section compares Tributary’s support for elastic services to two state-of-the-art resource managers designed for preemptible instances. Since neither system was designed for elastic services with latency SLOs, Tributary unsurprisingly performs significantly better.

Exosphere. We implemented ExoSphere’s allocation strategy, described in Sec. 2.2, with the following assumptions and modifications: (i) The ExoSphere paper did not specify whether the correlation between markets is recomputed as time moves on. In order to avoid the need to constantly reconstruct ExoSphere’s resource footprint, we assumed static correlation between markets. (ii) As the ExoSphere paper does not provide guidelines as to how to choose α , we experimented with a range of α from 1 to 10^9 . Higher α instructs ExoSphere to be more risk averse at the expense of higher cost.

Fig. 5 shows the normalized cost and percentage of “slow” requests served for Tributary and for ExoSphere with small (1) and large (10^9) values of α . These experiments were performed on a further scaled-up version of the *ClarkNet* trace (100x of already-scaled version), since ExoSphere was designed for 100s to 1000s of instances and performs poorly at a scale of 10s.⁴ In our experiments, we observed that Exosphere with a small α tends to acquire mainly the cheapest resources, inducing little diversity and increasing the number of “slow” requests in the event of preemptions. Tributary’s advantage in both cost and SLO attainment results from Tributary’s exploitation of spot instance characteristics (Sec. 5.3).

⁴At small scales, ExoSphere with low α had no resource diversity. With large α , it acquired too many resources, increasing its cost.

Proteus. We implemented Proteus’s allocation strategy, described in Sec. 2.2, modified to acquire only spot resources (reducing cost with no significant change in SLO attainment). Fig. 5 compares Tributary and Proteus for the *ClarkNet* trace, for two different scaling policies. While Proteus achieves lower cost than Tributary, it experiences a large increase in “slow” requests. This increase is due to Proteus not diversifying its resource pool, instead only acquiring resources based on reducing average per-core cost. When told by the scaling policy to acquire additional resources, similarly to AutoScale buffers (Sec. 5.3), Proteus is unable to match Tributary’s number of “slow” requests no matter how large the buffer (and, thus, how high the cost). This is once again due to the lack of diversity in the resources that Proteus acquires.

5.7 Prediction Model Evaluations

This section evaluates the accuracy of the preemption prediction models used by Tributary, which are described in Sec. 3.1. The recent Proteus system [16] used the historical median probability of preemption depending on the instance type, availability zone and the difference between the user bid price and the spot market price of the resource. Tributary improves prediction accuracy by using machine learning inference models trained with historical spot market data with engineered features. Fig. 6 shows the accuracy and F_1 scores for prediction models based on the historical median, a logistic regression classifier, a multilayer perceptron neural network (MLP NN) and a long short term memory recurrent neural network (LSTM RNN). These models were trained on spot market data from 06/06/16 – 01/22/17 and were evaluated on data from 01/23/17 – 03/20/17 for instance types *c4.large*, *c4.xlarge* and *c4.2xlarge* in *us-west-2*.

The output of the prediction models is whether the instance specified in a query will be preempted within the preemption window. Accuracy scores are calculated by the number of samples classified correctly divided by total number of samples. F_1 scores, which account for data skew, are a good accuracy measurement because the data set is skewed toward preemptions at lower *bid deltas* and non-preemptions at higher *bid deltas*. The LSTM RNN model provides the best accuracy and the best F_1 because it is able to capture the temporal nature of the AWS spot market. LSTM increases accuracy by 11% and the F_1 score by 27% compared to using the historical median. The MLP NN model performs worse than the historical median model for accuracy, but its F_1 score is higher because unlike the historical median model, the MLP model considers advanced features when predicting preemptions as described in Sec. 3.1. The increased accuracy of the LSTM RNN model translates to Tributary’s effectiveness. When using the LSTM RNN model, Tributary runs at $\approx 37\%$ less cost on the *ClarkNet* workload compared to Tributary using historical medians, because




Figure 6: Accuracies and F_1 scores (accounts for data skew) for predicting preemption of AWS spot instances. The LSTM RNN outperforms prior techniques (blue bar) by 11% on the accuracy metric and 27% on the F_1 score metric.

the historical median model overestimates the probability of preemption, causing Tributary to acquire more resources than necessary.

6 Conclusion

Tributary exploits AWS spot instances to meet latency SLOs for elastic services at lower cost. By predicting preemption probabilities and acquiring diverse resource footprints, Tributary can aggressively use collections of cheap spot instances to reliably meet SLOs even in the face of bulk preemptions. Our experiments show cost savings of 81–86% relative to using non-preemptible on-demand instances and 47–62% relative to traditional high-risk use of spot instances.

Tributary exploits AWS properties, such as dynamic spot markets and preemption based thereon. We believe its approach would also work for other clouds offering preemptible resources, if they expose enough information to predict preemption probabilities, which AWS provides via the visible spot market prices. Currently, *Google Cloud Engine* [5] does not expose such a signal for its preemptible instances. For private clouds, exposing preemption logs could provide the historical view, but even better predictions can be enabled by exposing scheduler state.

Acknowledgements: We thank our USENIX ATC18 reviewers and our shepherd, Christopher Stewart, for valuable suggestions. We thank Henggang Cui for his feedback. We thank the members and companies of the PDL Consortium (Alibaba, Broadcom, Dell EMC, Facebook, Google, HPE, Hitachi, IBM, Intel, Micron, Microsoft, MongoDB, NetApp, Oracle, Salesforce, Samsung, Seagate, Two Sigma, Toshiba, Veritas and Western Digital) for their interest, insights, feedback, and support. This research is supported in part by Intel as part of the Intel STC for Visual Cloud Systems (ISTC-VCS), National Science Foundation under awards CNS-1042537, CCF-1533858, CNS-1042543 (PRObE [13]), and DARPA Grant FA87501220324. Lastly, one of the co-authors is supported in part by DHS Award HSHQDC-16-3-00083, NSF CISE Expeditions Award CCF-1139158, and gifts from Alibaba, Amazon Web Services, Ant Financial, CapitalOne, Ericsson, GE, Google, Huawei, Intel, IBM, Microsoft, Scotiabank, Splunk and VMware.

References

- [1] Amazon EC2 Spot Instances. <https://aws.amazon.com/ec2/spot>.
- [2] AWS Autoscale. <https://aws.amazon.com/autoscaling/>.
- [3] AWS EC2. <http://aws.amazon.com/ec2/>.
- [4] AWS EC2 Spot Fleet. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-fleet.html>.
- [5] Google Compute Engine. <https://cloud.google.com/compute/>.
- [6] Spot Bid Advisor. <https://aws.amazon.com/ec2/spot/bid-advisor/>.
- [7] Tensorflow serving. <https://tensorflow.github.io/serving>.
- [8] ABADI, M., BARHAM, P., CHEN, J., CHEN, Z., DAVIS, A., DEAN, J., DEVIN, M., GHEMAWAT, S., IRVING, G., ISARD, M., KUDLUR, M., LEVENBERG, J., MONGA, R., MOORE, S., MURRAY, D. G., STEINER, B., TUCKER, P., VASUDEVAN, V., WARDEN, P., WICKE, M., YU, Y., AND ZHENG, X. Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation (OSDI 16)*.
- [9] AGMON BEN-YEHUDA, O., BEN-YEHUDA, M., SCHUSTER, A., AND TSAFRIR, D. Deconstructing Amazon EC2 spot instance pricing. *ACM Transactions on Economics and Computation* 1, 3 (2013), 16.
- [10] DANZIG, P., MOGUL, J., PAXSON, V., AND SCHWARTZ, M. The internet traffic archive. URL: <http://ita.ee.lbl.gov/> (2000).
- [11] GALANTE, G., AND BONA, L. C. E. D. A survey on cloud computing elasticity. In *Proceedings of the 2012 IEEE/ACM Fifth International Conference on Utility and Cloud Computing* (Washington, DC, USA, 2012), UCC '12, IEEE Computer Society, pp. 263–270.
- [12] GANDHI, A., HARCHOL-BALTER, M., RAGHUNATHAN, R., AND KOZUCH, M. A. Autoscale: Dynamic, robust capacity management for multi-tier data centers. *ACM Trans. Comput. Syst.* 30, 4 (Nov. 2012), 14:1–14:26.
- [13] GIBSON, G., GRIDER, G., JACOBSON, A., AND LLOYD, W. Probe: A thousand-node experimental cluster for computer systems research. *USENIX* 38, 3 (2013).
- [14] GONG, Z., GU, X., AND WILKES, J. Press: Predictive elastic resource scaling for cloud systems. In *6th IEEE/IFIP International Conference on Network and Service Management (CNSM 2010)* (Niagara Falls, Canada, 2010).
- [15] GROUP, W. N. R., ET AL. Wits: Waikato internet traffic storage. URL: <http://wand.net.nz/wits/index.php>.
- [16] HARLAP, A., TUMANOV, A., CHUNG, A., GANGER, G. R., AND GIBBONS, P. B. Proteus: Agile ML elasticity through tiered reliability in dynamic resource markets. In *Proceedings of the Twelfth European Conference on Computer Systems* (New York, NY, USA, 2017), EuroSys '17, ACM, pp. 589–604.
- [17] HOXMEIER, J. A., AND DICESARE, C. System response time and user satisfaction: An experimental study of browser-based applications. *AMCIS 2000 Proceedings* (2000), 347.
- [18] KALISKI, B. Pkcs# 5: Password-based cryptography specification version 2.0.
- [19] KOHAVI, R., AND LONGBOTHAM, R. Online experiments: Lessons learned. *Computer* 40, 9 (2007).
- [20] LIU, H., AND WEE, S. *Web Server Farm in the Cloud: Performance Evaluation and Dynamic Architecture*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 369–380.
- [21] MARATHE, A., HARRIS, R., LOWENTHAL, D., DE SUPINSKI, B. R., ROUNTREE, B., AND SCHULZ, M. Exploiting redundancy for cost-effective, time-constrained execution of HPC applications on Amazon EC2. In *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing* (2014), ACM, pp. 279–290.
- [22] MULLER, K., AND VIGNAUX, T. Simpy: Simulating systems in python. *ONLamp. com Python Devcenter* (2003).
- [23] SHARMA, P., GUO, T., HE, X., IRWIN, D., AND SHENOY, P. Flint: Batch-interactive data-intensive processing on transient servers. In *Proceedings of the 11th European Conference on Computer Systems (EuroSys 16)* (2016), ACM, p. 6.
- [24] SHARMA, P., IRWIN, D., AND SHENOY, P. Portfolio-driven resource management for transient cloud servers. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1, 1 (2017), 5.
- [25] SHARMA, P., LEE, S., GUO, T., IRWIN, D., AND SHENOY, P. Spotcheck: Designing a derivative iaas cloud on the spot market. In *Proceedings of the Tenth European Conference on Computer Systems* (2015), ACM, p. 16.
- [26] SHARMA, U., SHENOY, P., SAHU, S., AND SHAIKH, A. A cost-aware elasticity provisioning system for the cloud. In *2011 31st International Conference on Distributed Computing Systems* (June 2011), pp. 559–570.
- [27] SHASTRI, S., AND IRWIN, D. Hotspot: Automated server hopping in cloud spot markets.
- [28] SOUDERS, S. Velocity and the bottom line. In *Velocity (Web Performance and Operations Conference)* (2009).
- [29] SUNDERMEYER, M., SCHLÜTER, R., AND NEY, H. LSTM neural networks for language modeling. In *Interspeech* (2012), pp. 194–197.
- [30] TANG, S., YUAN, J., AND LI, X.-Y. Towards optimal bidding strategy for Amazon EC2 cloud spot instance. In *Proceedings of the 5th IEEE International Conference on Cloud Computing (CLOUD 12)* (2012), IEEE, pp. 91–98.
- [31] WANG, W., LI, B., AND LIANG, B. To reserve or not to reserve: Optimal online multi-instance acquisition in iaas clouds. In *ICAC* (2013), pp. 13–22.
- [32] XU, Z., STEWART, C., DENG, N., AND WANG, X. Blending on-demand and spot instances to lower costs for in-memory storage. In *INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications, IEEE* (2016), IEEE, pp. 1–9.
- [33] ZHENG, L., JOE-WONG, C., TAN, C. W., CHIANG, M., AND WANG, X. How to bid the cloud. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (2015), ACM, pp. 71–84.