

## About Us

Exotanium, a cloud resource optimization and management platform that uses cutting-edge distributed container infrastructure to help enterprises use cloud resources more efficiently and reduce cloud compute costs by up to 90 percent. The company's technology comes from 10+ years of research done at Cornell University. Exotanium is dedicated to developing cloud computing technology that empowers enterprises to better control their cloud computing resources while enhancing both operational security and performance.


## Our Solution

### PROBLEM

- ❖ **Millions of dollars spent** for long running workloads
- ❖ **Over provisioned** and **idle Virtual** machines
- ❖ **Underutilization** of cloud resources

**SOLUTION:** **xspot** 

**xSpot** provides high reliability while reducing cost up to 2x - 4x for **stateful and stateless workloads**



### CASE STUDY

Company: Fortune 1000 Computer Aided Design Company

#### The Workload:

- ❖ Computer aided design analysis
- ❖ Stateful, legacy, and long running

#### Results with X-Spot

- ❖ **>60%** cost saving
- ❖ **<6%** performance overhead



#### Significant Cost Reduction

Up to 90% reduction in costs due to the use of cloud spot instances.



#### Patented Technology

Patented technology provides same availability as regular instances.



#### Seamless without Interruption

Uses Live Transparent Migration, without interruption

❖ Container size: 32 CPU, 128 GB memory

❖ Migrate transparently without affecting the front-end

## Underlying Mechanism

### Have your cake and eat it: Reliably running stateful application containers in cheap spot markets

Virtual Machine “spot markets” provide virtual machine resources often at a fraction (10-20%) of the price of regular cloud offering. Major cloud providers such as Amazon AWS, Microsoft Azure, and Google Cloud allow you to bid for unused resources. Amazon calls them *Spot Instances*, Microsoft calls them *Low-Priority VMs*, and Google calls them *Preemptible VM Instances*. The only catch is, the providers reserve the right to reclaim the instances at any time. This makes a spot market useful mostly for *short-running* and *stateless* tasks that can be quickly restarted on a different VM if need be. A good example is a media processing task such as rendering.

Can we make the spot market useful to arbitrary containerized applications? This gives rise to various challenges. For example, many container applications are stateful. Particularly containerized legacy applications that were not designed to run in the cloud cannot simply be rebooted at will, as each such reboot would result in an unacceptably long interruption of service time. Are those spot markets therefore unusable for stateful applications?

A promising technology that might make spot markets useful to such applications is *Live Virtual Machine Migration*. Cloud providers internally use live virtual machine migration all the time to balance the load of VMs across their compute clusters. If made available to cloud users, they could migrate containers into the spot market when there are plenty of unused resources and the spot market is reliable and move them out of the spot market when resource reclamation is likely.

VM migration works in two phases. In the *pre-copy phase*, the memory image of the VM is copied from the source host to the destination host *while it is running*. In the *switch phase*, the VM on the source host is first paused. Any memory pages that were modified after they were copied, hopefully not too many, are then copied again. Also, some routing tables need to be updated so messages are correctly routed to the new host. Next, the VM is restarted on the destination VM. Finally, the VM on the source host is terminated. While the whole process of copying might take a minute or two, most of that time is spent in the pre-copy phase while the VM is running. The switch phase typically only takes a couple seconds, after which the VM is executing on the new host.

There are, however, various obstacles for users to leverage this technology. First, cloud providers do not offer virtual machine migration between regular cloud resources and spot market resources, so how can we make it available to end-users? Second, how do we make

migration transparent to container runtime systems such as Docker and Kubernetes, particularly if the container is stateful? Third, how would one decide when it is safe to run a container in the spot market and when not?

The first problem can be overcome with so-called *nested virtualization*. Instead of running an operating system such as Linux or Windows directly inside a virtual machine, one can, in theory, run a virtual machine monitor inside a virtual machine. The nested virtual machine monitor would be under the control of the user, and users can then leverage existing live VM migration tools to move a virtual machine to another virtual machine monitor under their control. Nested virtualization can even support other cost-saving features. For example, a user could securely run multiple applications inside the same cloud VM when load is low and migrate them to different cloud VMs when the load is high. If not designed carefully, however, nested virtualization could result in significant performance overheads. The paper "[Cloud Mobility for Geographically Shifting Workloads](#)" by Ying Xiong and Hakim Weatherspoon in *The New Stack* (July 24 2020) provides more details on nested virtualization and various interesting use cases for legacy applications.

Even when virtual machine migration is available to end-users, the second problem to solve is to make it efficient and transparent to container runtime systems. The Cornell paper titled "X-Containers: Breaking Down Barriers to Improve Performance and Isolation of Cloud-Native Containers" by Zhiming Shen et al. (Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems) shows how to build a 100% Linux-compatible container framework that supports live-migratable and efficient containers. But even then you would still have to implement the Open Container Initiative (OCI) and Kubernetes Container Runtime Interface (CRI) to be able to use the technology with container frameworks..

To overcome the third problem, one could consider the termination warnings that some spot markets provide. However, these typically do not afford enough time to migrate the application out of the spot market. A better option is to proactively migrate out of the spot market to regular cloud offerings. Such decisions would be based on monitoring the spot market and using heuristics or machine learning to decide when premature termination is likely. The paper "[Smart Spot Instances for the Supercloud](#)" by Qin Jia, Zhiming Shen, Weijia Song, Robbert van Renesse, and Hakim Weatherspoon in the CrossCloud Workshop, April 2016 investigates various heuristics.

These concepts originated in research at Cornell University and are now under development at a company called *Exotanium*. Exotanium provides a nested virtualization platform that can run on any cloud platform, homogenizing the diverse cloud infrastructures offered today. One VM image can now run in any cloud and can be live-migrated between availability zones and even across different cloud providers. Further, Exotanium provides an OCI and CRI compatible application container runtime. Exotanium further provides cost-saving products that leverage the spot market and the ability to securely consolidate underloaded containers.

## Join Our Team

### **Solution Architects**

Exotanium is searching for Solution Architects to integrate our next-generation cloud-native container management platform with customer applications.

### **Software Engineer in Tests**

Exotanium is searching for Software Engineers in Test to build test automation for our next-generation cloud-native container management platform with customer applications.

### **Software Engineer - Cloud Integration**

Exotanium is searching for Cloud Integration Software Engineers to build out our next-generation cloud-native container management platform.

### **Software Engineer - Container Runtime**

Exotanium is searching for Container Runtime Software Engineers to build out our next-generation cloud-native container management platform.

### **Software Engineer (Linux Kernel and Hypervisor Development)**

Exotanium, Inc. Software Engineer (Linux Kernel and Hypervisor Development) to work with a team of developers to extend our platform's capabilities to new cloud services.

### **Internships**

Exotanium is seeking full-time Software Engineers and Research/Development Interns to join our team. Internship positions are not listed specifically and will be determined case by case.