

# The XenoServer Open Platform

## Deploying global-scale services for fun and profit

Evangelos Kotsovinos and David Spence

### Summary

- The XenoServer project is building a global-scale, public Infrastructure for distributed computing
- **XenoServers** can host and securely execute untrusted and unsafe code in exchange for money
- **XenoCorp** acts as a trusted third party –much like VISA for card purchases–
- The **XenoServer Information Service (XIS)** operates as the “yellow pages” of the platform, storing the specifications of active XenoServers
- **XenoStore** provides unified, flexible distributed storage for the participants in the platform
- **XenoSearch** performs complex search operations on XIS on behalf of the clients

### Research challenges

- Distributed platform for code deployment
- Safe execution of any untrusted task, without the need for “safe” or “signed” code
- Efficient and flexible XenoServer discovery
- Providing QoS guarantees
- Charging & auditing mechanisms
- Unified adaptable storage
- Distributed trust management
- Viable self-financing economic model

### Features

- Ease of deployment of global-scale services
- Low cost of entry
  - Monetary
  - Effort
- Incrementally scalable on-demand
- Short timescales
- Fine-grained charging
- No need for trusted or reliable code
- Security and QoS isolation
- Logging and Forensic Auditing
- No restrictions on O/S or programming APIs
- Tackles all related problems at the same time

### References

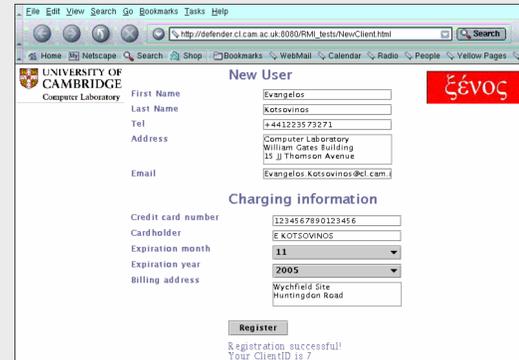
- **Controlling the XenoServer Open Platform.** Steven Hand, Tim Harris, Evangelos Kotsovinos and Ian Pratt. *In Proc. of 6th Intl. Conference on Open Architectures and Network Programming (IEEE OPENARCH '03)*
- **Xen 2002.** P.R. Barham, B. Dragovic, K.A. Fraser, S.M. Hand, T.L. Harris, A.C. Ho, E. Kotsovinos, A.V.S. Madhavapeddy, R. Neugebauer, I.A. Pratt, A.K. Warfield. *Computer Lab Technical Report 553, Jan 2003*
- **XenoSearch: distributed resource discovery in the XenoServer Open Platform,** David Spence, Tim Harris, *12th IEEE Intl Symp. on High Perf. Dist. Computing (HPDC-12)*
- **XenoTrust: Event-based distributed trust management.** Boris Dragovic, Evangelos Kotsovinos, Steven Hand and Peter Pletzuch. *In Proc. of the 2nd IEEE Intl Workshop on Trust and Privacy in Digital Business (DEXA-TrustBus'03)*
- **Managing Trust and Reputation in the XenoServer Open Platform** Boris Dragovic, Steven Hand, Tim Harris, Evangelos Kotsovinos and Andrew Twigg. *To appear in Proc. of the 1st Intl Conference on Trust Management, May 2003, Heraklion, Crete.*

<http://www.xenoserver.org>

## 1 Register

XenoServers and clients register with XenoCorp in order to obtain the **authentication credentials** that allow participation in the Platform. Registration binds an entity to a **real-world identity**.

Before using resources on XenoServers, clients need to create **purchase orders** with XenoCorp. This represents a commitment to funding sessions, subject to certain constraints. At that point, XenoCorp can check the credit-worthiness of the client as well as ring-fence the portion issued as a purchase order.



## 3 Create sessions

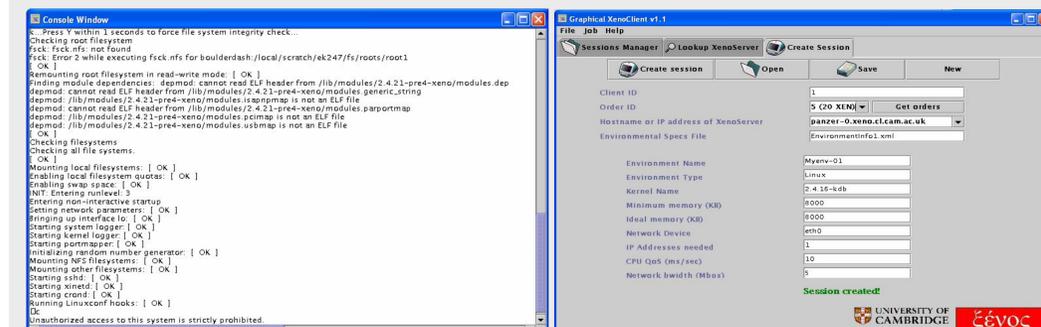
After locating and choosing the XenoServer that will be used to host his task, the client requests a session creation on the XenoServer, providing the specifications and requirements for his task. The creation of a session follows a positive admission control decision, and establishes an **agreement** between the client and the XenoServer regarding the **resources** to be provided, and the **payment** to be made. A **new execution environment** (instance of an operating system such as Linux, BSD or WinXP, or an execution platform such as JVM) is launched for the client.

The agreement established at session creation specifies that both the environmental (such as type of OS and kernel version) and QoS (such as milliseconds of CPU time per wall-clock second or Mbps of network bandwidth) requirements of the task will be met by the XenoServer, as well as the pricing scheme that will be used.

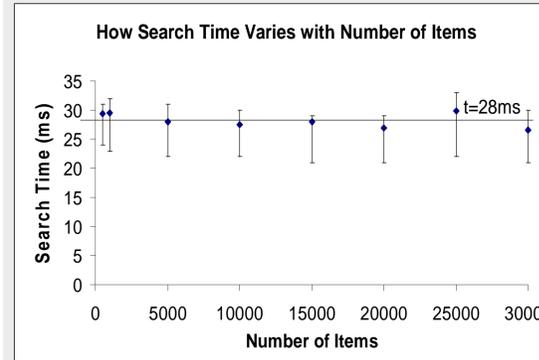
Our prototype XenoServer is based on a low-level component, termed the **Xen hypervisor**, which virtualizes the physical resources of the machine, apportioning them between the various environments that

it hosts by creating a **virtual machine** for each one. Protection and resource isolation happen at the lowest possible level, making sure that unsafe and unverified tasks can only be mischievous inside their own environment.

Higher-level decisions, such as admission control, as well as the association of sessions with environments in a one-to-one or one-to-many fashion, are taken by the control-plane software, running in the administrative environment called domain0.



## 2 Search



Clients **discover** the most **suitable XenoServers** to deploy their code on by querying XenoSearch. Queries can specify the environment required (e.g. type & version of operating system) and various QoS attributes (e.g. amount of memory or ms of CPU per wall-clock second). Complex queries, such as requesting XenoServers that are close to a set of hosts, are represented easily.

XenoSearch indexes **advertisements** periodically produced by XenoServers. An advertisement contains information on the **identity, ownership** and **location** of the XenoServer, its total and available **resources** and their **prices**. XenoSearch converts advertisements to points in an n-dimensional space, different dimensions representing different attributes including topological location and QoS attributes. XenoSearch uses XenoStore, which provides a unified distributed storage service for the XenoServer platform, to store, distribute and cache advertisements.

XenoSearch employs a novel scheme to **summarize** the advertisements efficiently within a range of a dimension. Each dimension is stored in a separate Distributed Hash Table, these are summarized, and the resulting summaries are also stored within the DHT. Ranges can then be queried and combined to obtain a summary of an arbitrary range. Complex searches are built up from ranges and the union and intersection operators. Finally, the advertisements in the summary are queried to find the exact set that matches the query.

## 4 Deploy

Once a session is created, the client can connect to the new environment he has leased to start **deploying tasks**. Deployment can be done in any way the client is familiar with –for instance, using ssh or a remote method invocation platform. **Ongoing management** of environments, such as dynamic allocation and deallocation of resources, is supported.

The XenoServer will **account** for the resources consumed by each of the execution environments, associate them with the respective purchase orders and **claim payments** from XenoCorp. Tasks' activities are being logged for spotting potential cases of unlawful behaviour. Performance evaluation experiments we have carried out show that the overhead imposed by Xen's resource virtualization is very limited, which means that clients experience no significant performance drop when running their tasks on XenoServer environments instead of commodity operating systems. Our hypervisor significantly outperforms similar commercial products, such as the VMware ESX server.

