Barriers to Cloud Adoption and Research Opportunities

Prashant Shenoy

University of Massachusetts Amherst

Cloud Economics

- Cloud platforms: lease hardware and software rather than own
    - Argument: Leasing is cheaper than owning an IT infrastructure
- Large enterprises: already outsource some/all of their IT infrastructure
    - Services company takes over/manages IT systems
    - Costs benefits of “leasing” are already built into the outsourcing model
- Cloud platforms: on-demand resource provisioning
    - Looking to provide private cloud services as part of their offerings
    - **Challenge**: How to design a economical private cloud platforms?
        - Potentially lower multiplexing benefits than public clouds

Manageability and the Cloud

- Large enterprises : IT infrastructure spread over multiple data centers
    - Typically few “large” and many “small” data centers
- Use of Enterprise data center management tools
    - IBM Tivoli, HP Openview ....
- Cloud is just another virtualized “data center” with resources
    - Need to treat cloud resources similar to local IT resources
    - Cloud resources should integrate “seamlessly” into the IT infrastructure
- **Challenge**: Enterprise Management tools need to evolve to handle cloud resources

Network Resources and the Cloud

- Cloud platforms offer compute / storage / software stacks on-demand
- Cloud boundary typically ends to the network “edge”
- Enterprises need control over network (in addition to compute/storage)
  - VPNs, bandwidth provisioning between data centers and “branch offices”
  - Network reconfiguration needs to be made on-demand
- Challenge: Extend boundary of cloud platform to allow network control
  - Cloud tools need to work well with network management tools

Barriers to Cloud Adoption and Research Opportunities

Ion Stoica

UC Berkeley

Two Challenges (among Many, Many Others…)

- Provide meaningful SLAs
  - Availability, Reliability, Performance, Predictability
  - "One throat to choke"

- Data security
  - Privacy, access control & protection, audition

Predictable Performance

- Holly grail: give the application the illusion it runs on a dedicated cluster
- Research opportunities:
  - Achieve high utilization while supporting interactive applications
  - Isolation across resources and services
    - Caches, I/O, and Networking most challenging
  - Dynamically scale up and down in response to load to maintain performance target

Multi-Datacenter Support

- Usage scenarios:
  - Hybrid use of private and public clouds
  - Replication to improve availability
  - Partitioning to improve scalability
- Research opportunities:
  - Data consistency
  - Data-center and service instance selection
  - Unified APIs across cloud providers

Data Security

- Goals
  - Provide privacy, confidentiality, deletion, non-repudiation, auditability, while…
  - … allowing services to run across multiple administrative domains
- Research opportunities:
  - Privacy-preserving queries
  - Queries on encrypted data
  - Leverage TCM for OS/software platforms

Systems and Internet Infrastructure Security

Network and Security Research Center
Department of Computer Science and Engineering
Pennsylvania State University, University Park PA

Cloud Security: Challenges and Opportunities

Trent Jaeger w/ Joshua Schiffman
USENIX HotCloud Workshop
June 22, 2010

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

Security View

Cloud Data Center

Your Data

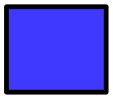

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

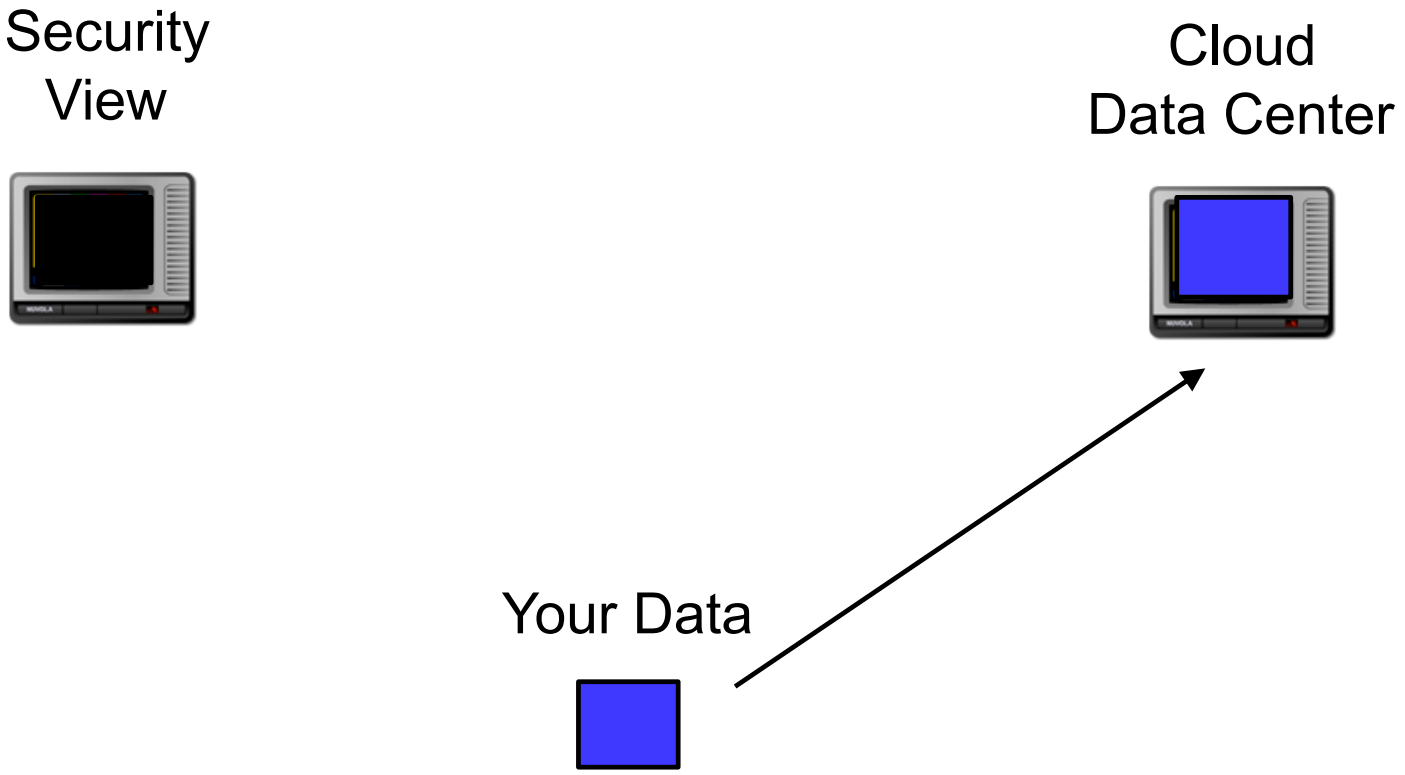

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

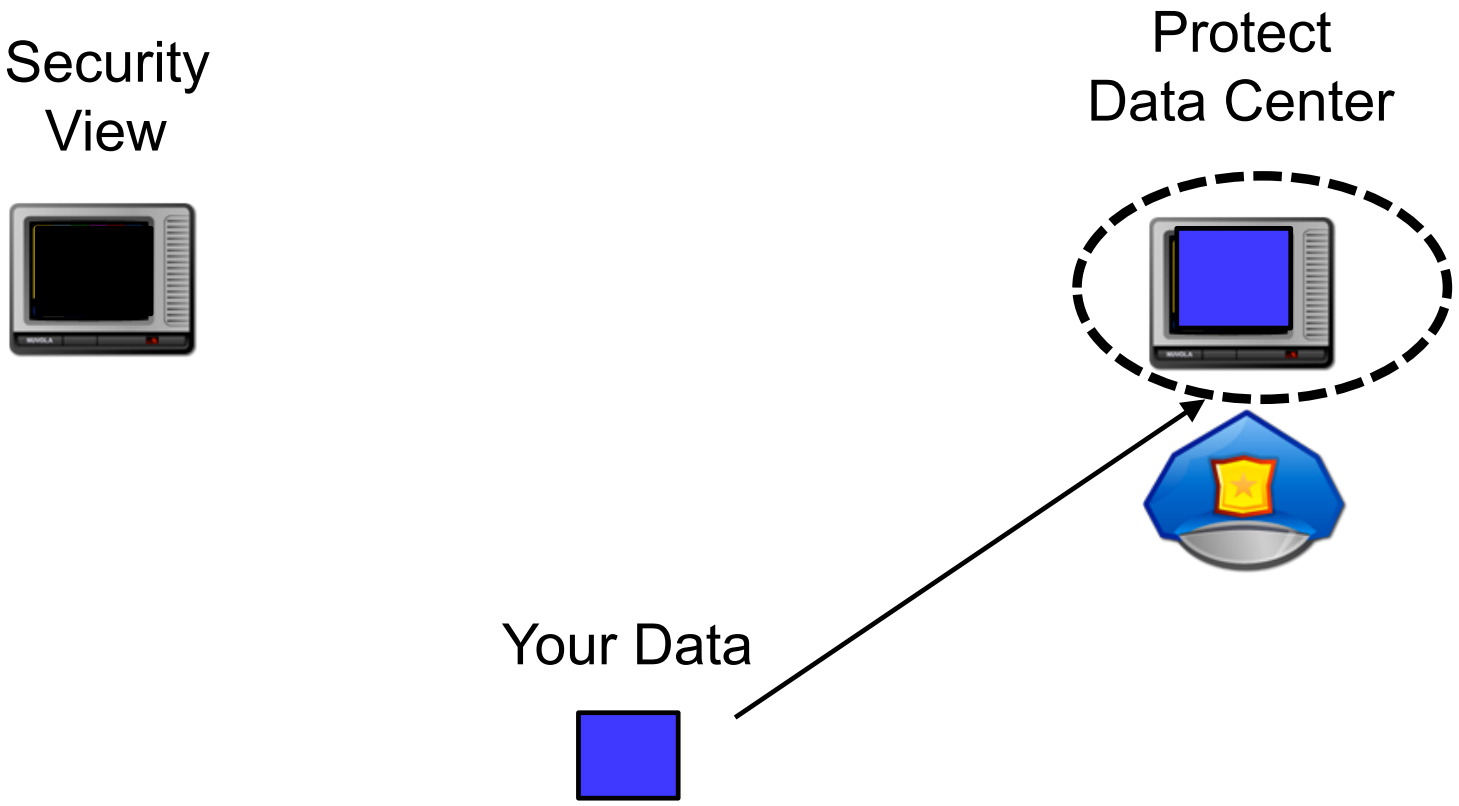

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

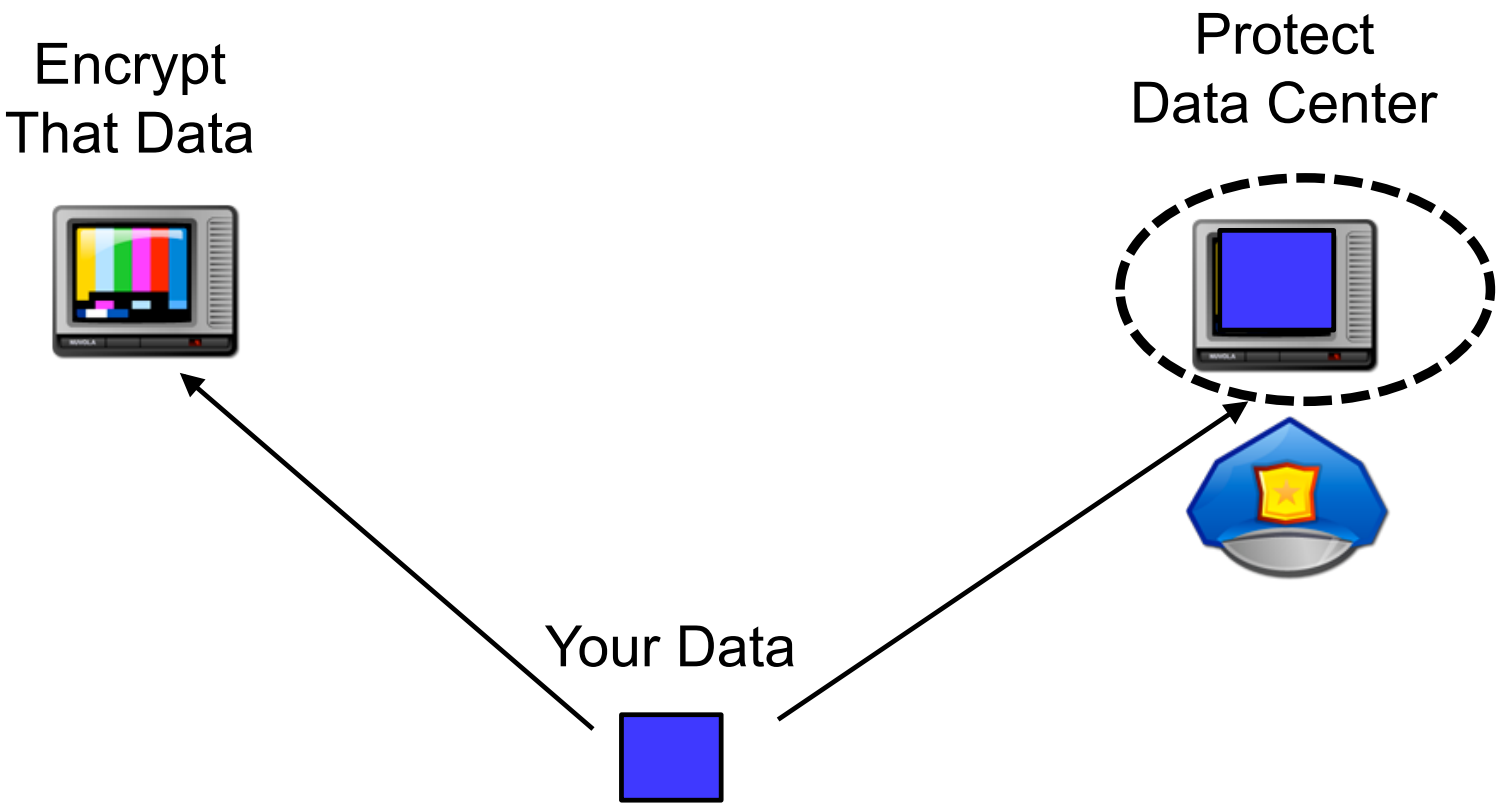

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

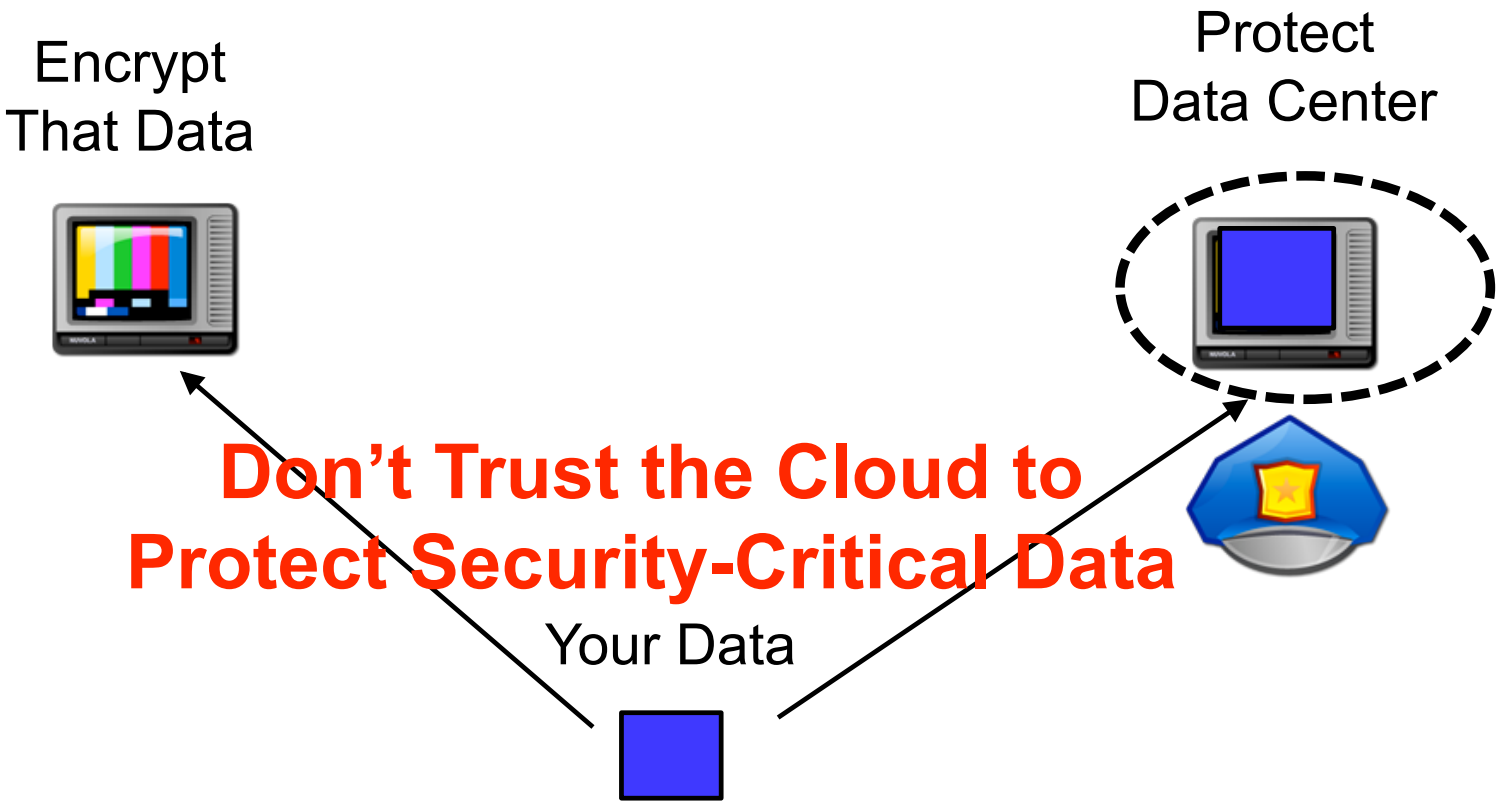

Cloud: Can It Help?

- I think (hope) so
  - Should enable scalable administration
    - Could create world-class node controllers
    - Help customers choose among cloud instance configurations
    - Help customers configure their firewall/VLAN
    - From help to risk mitigation/minimization
- But, there are challenges
  - Host security has a number of problems
  - Insiders are outsiders to you
  - Covert channels

Trustworthy Base

- *Challenge*: Verifying the initial state of the entire attestation framework
- Potentially large TCB to verify
  - Code and data
  - Need methods for assessing dynamic data
- Provenance of system to a trusted origin
  - Root of Trust for Installation [ACSAC '07]
  - To Deploy a VM [ACSAC '09]
- Alternative is to assess the impact of the data

Integrity Monitoring

- *Challenge*: Prove correct execution of all integrity-relevant events
- Record events for later verification
  - Verify after the fact
  - Difficult to evaluate without context
- Enforcement can reduce verification effort
  - Must verify enforcement mechanism and policy
  - Allow untrusted data to VM, process, and program
- Can measure this across VMs, between OS/process

Cloud Security Perceptions

- Security: What do we need to do to protect a customer's data running in the cloud?

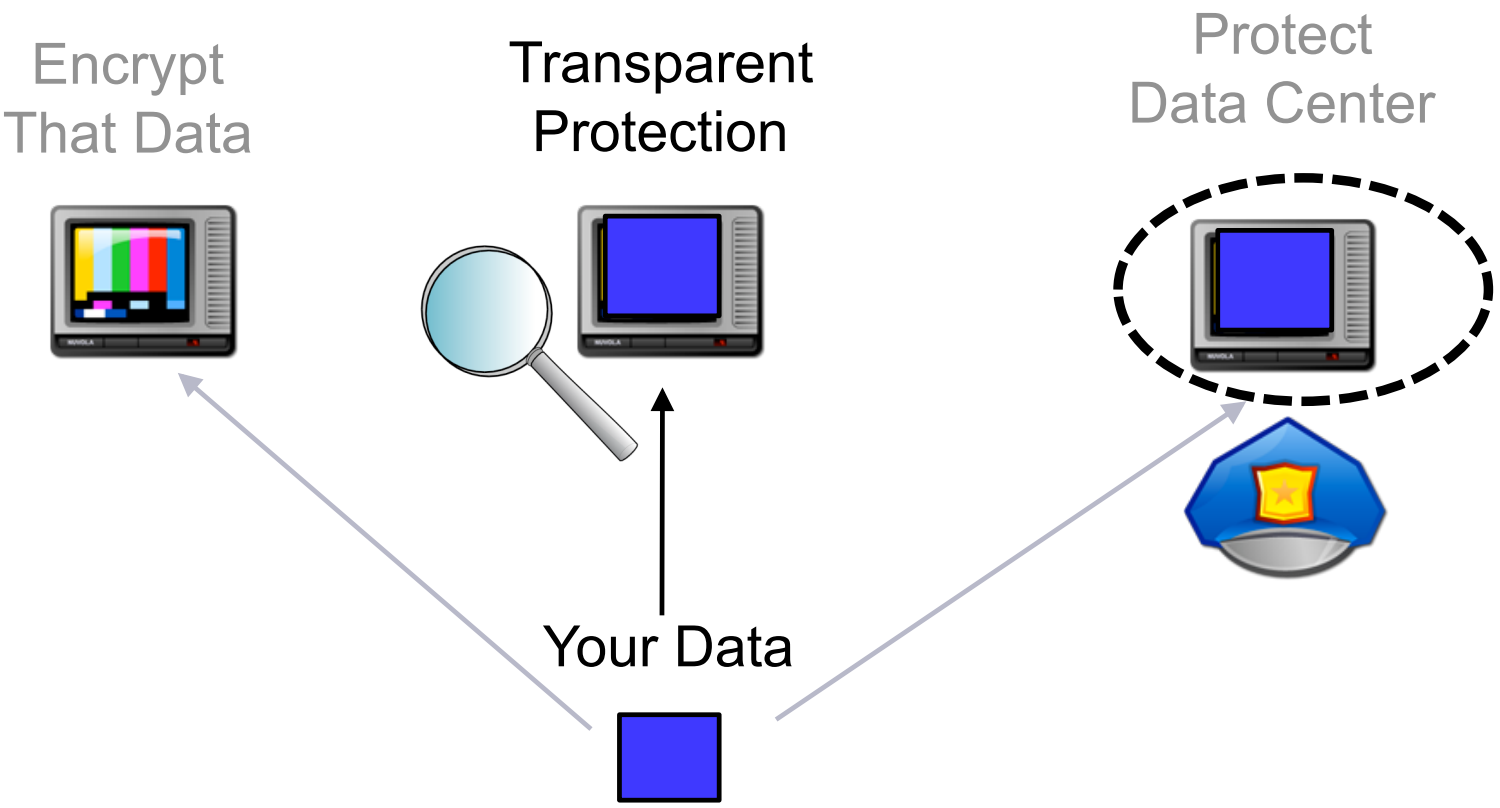

Verifying Cloud Nodes

- Servers have a greater incentive to prove compliance
  - Proof of the node's correctness
  - Supplement SSL Certificates
  - Large companies can manage internal PKI
- Use hardware attestation to prove compliance
- Adoption challenges
  - Performance
  - Privacy concerns
  - Don't want to be restricted by complicated processes

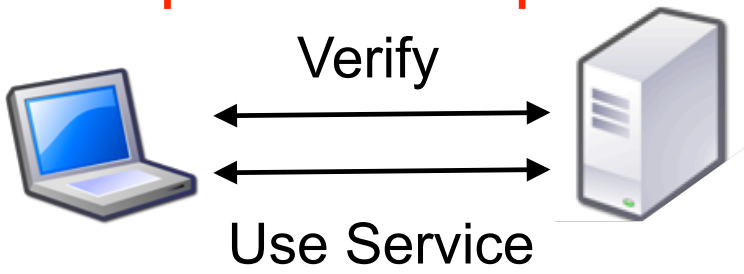

Summary

- Cloud is not trusted with security-critical data
- How should we ensure that security-critical data is not leaked?
  - And prove it to customers…
- Armed guards are not sufficient
  - But, computing with encrypted data is not what we had in mind
- Claim: Build systems to preserve integrity and prove via attestation

Thank you

Trent Jaeger (tjaeger@cse.psu.edu)

http://www.trentjaeger.com/

SIIS Laboratory (http://siis.cse.psu.edu)

Cloud Computing

Orran Krieger
Cloud Architect

© 2009 VMware Inc. All rights reserved

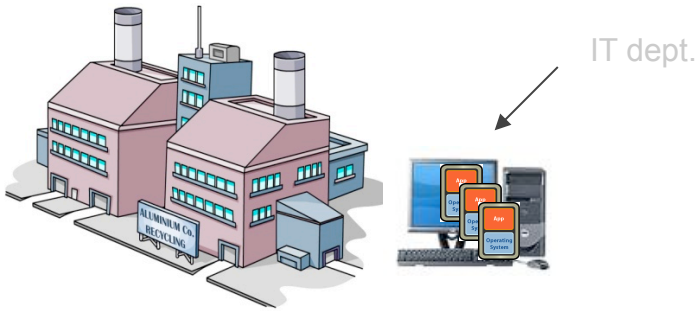

Joe’s Widget Co.

vmware®

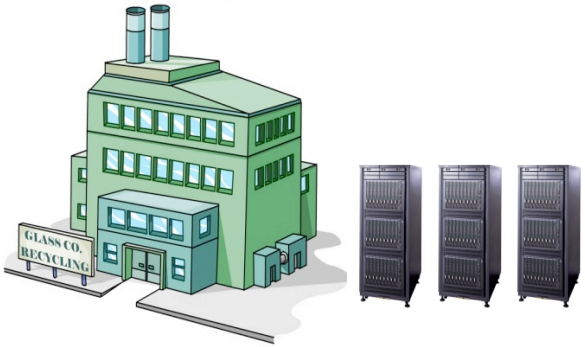

Host it R us.

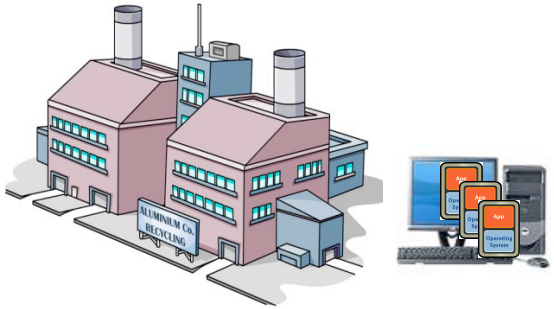

Joe's Widget Co.

vmware®

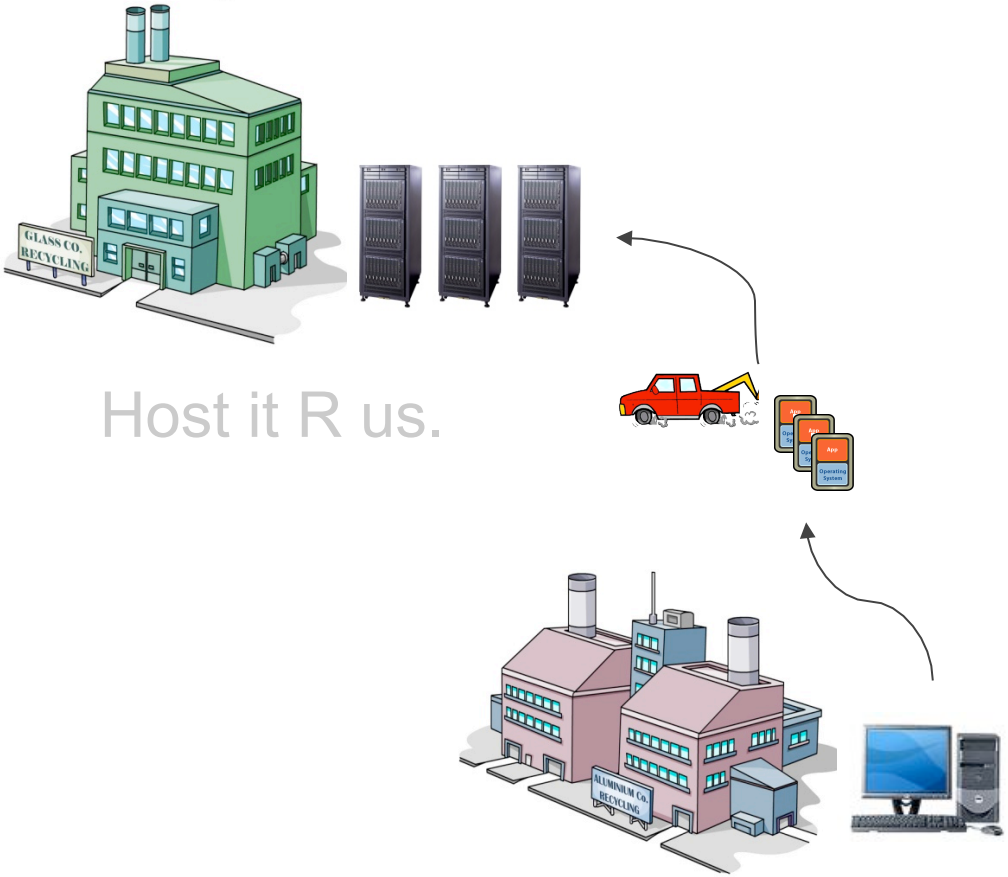
GLASS CO.
RECYCLING
Host it R us.
ALUMINUM Co.
RECYCLING
Joe's Widget Co.

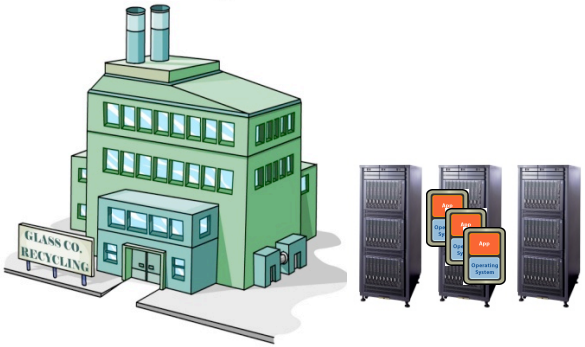

Host it R us.

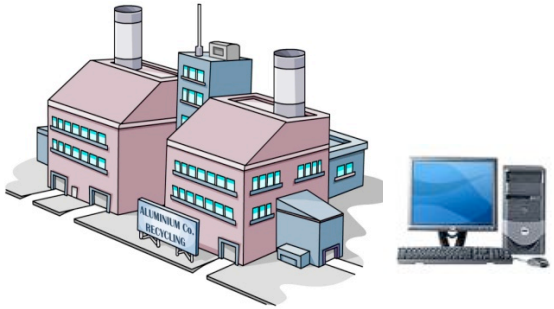

Joe’s Widget Co.

vmware®

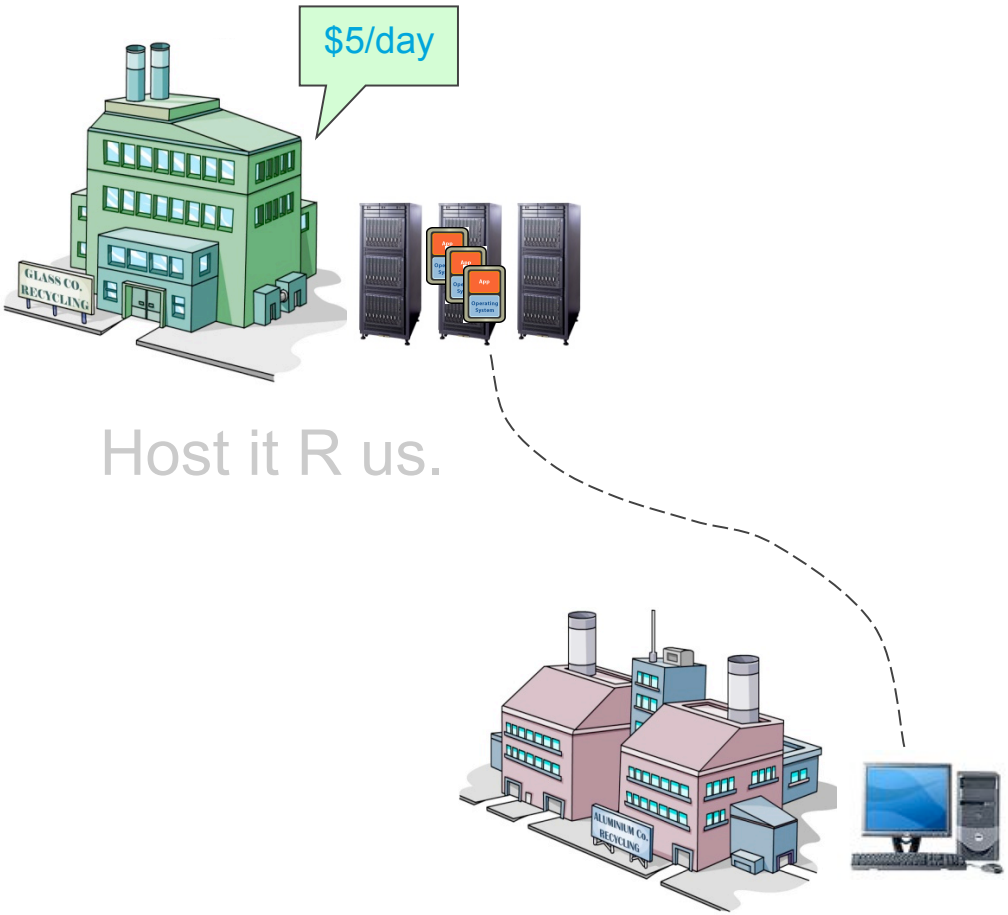
$5/day
GLASS CO.
RECYCLING
App
Operating System
Host it R us.
ALUMINUM Co.
RECYCLING
Joe's Widget Co.

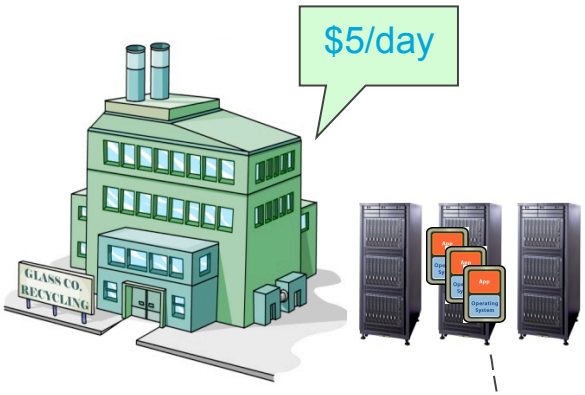

Host it R us.

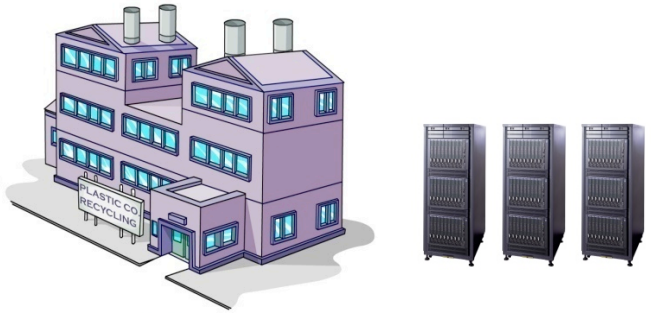

Host 4 Less

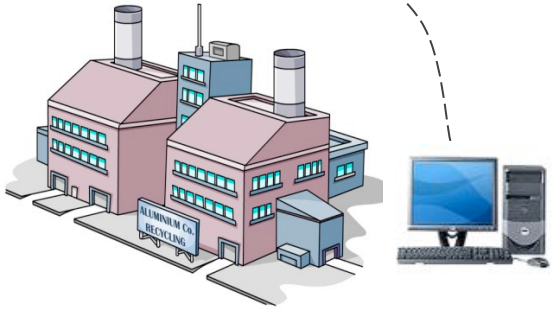

Joe's Widget Co.

vmware®

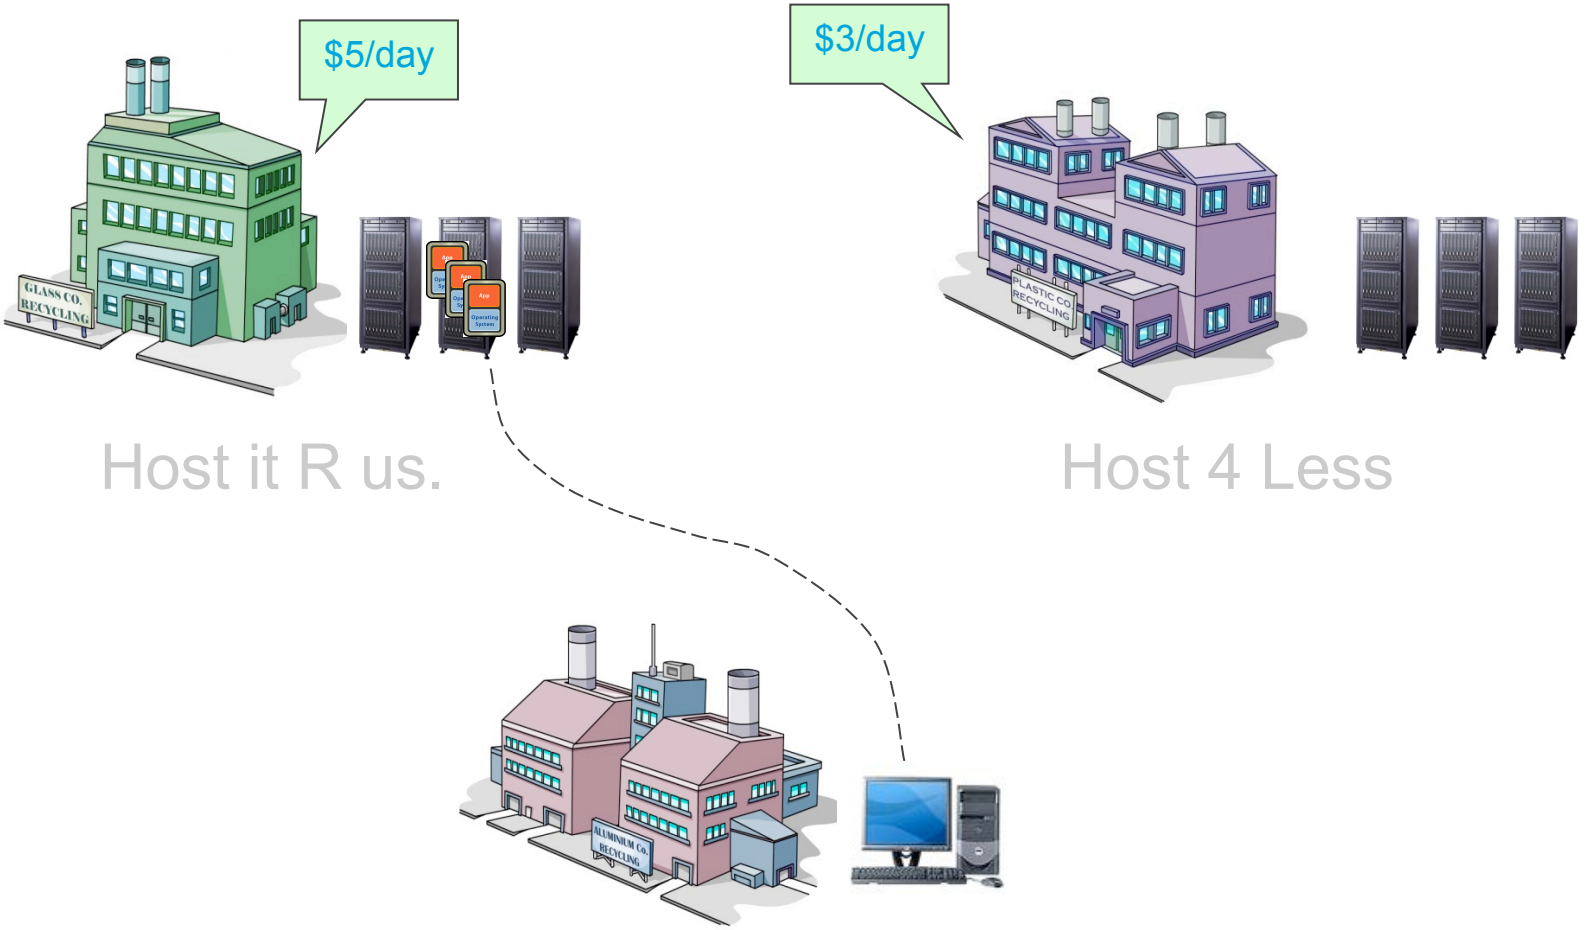
$5/day
$3/day
GLASS CO.
RECYCLING
App
Operating System
Host it R us.
Host 4 Less
ALUMINUM Co.
RECYCLING
Joe's Widget Co.

vmware®

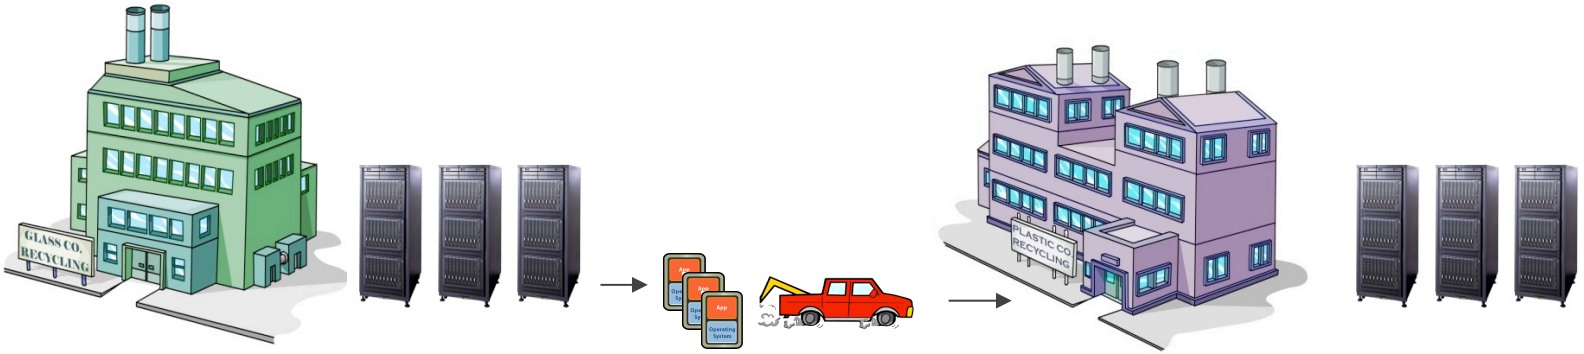

Host it R us.

Host 4 Less

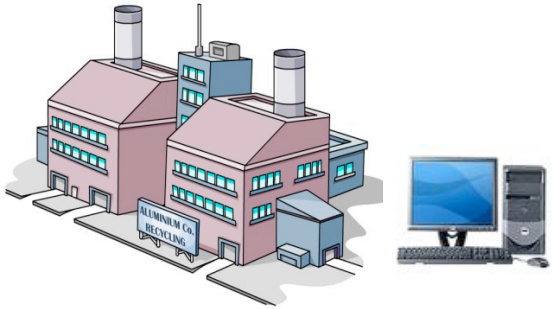

Joe’s Widget Co.

vmware®

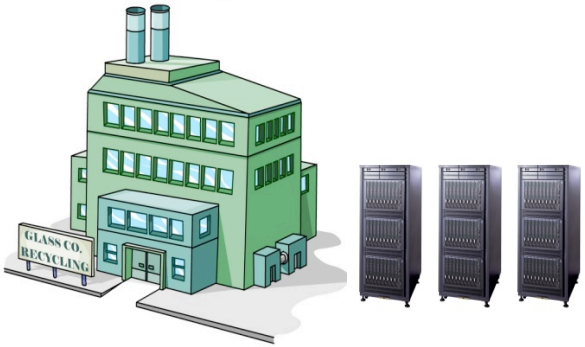

Host it R us.

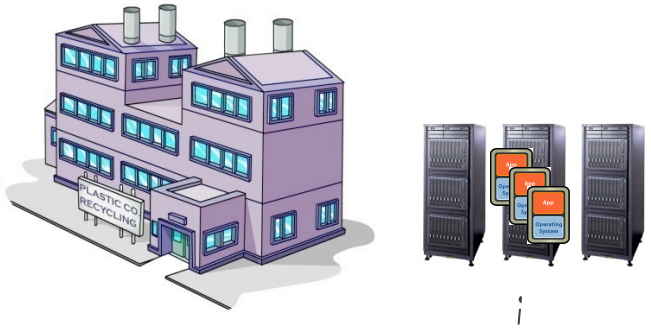

Host 4 Less

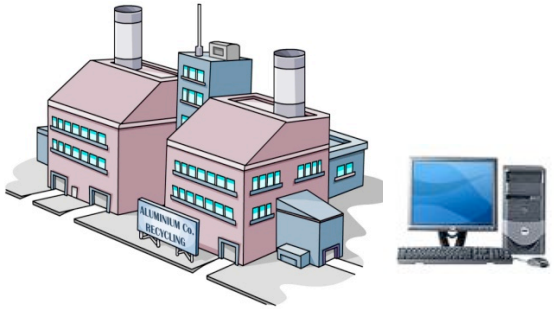

Joe's Widget Co.

vmware®

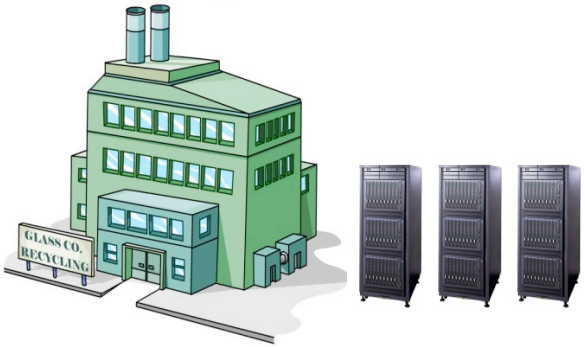

Host it R us.

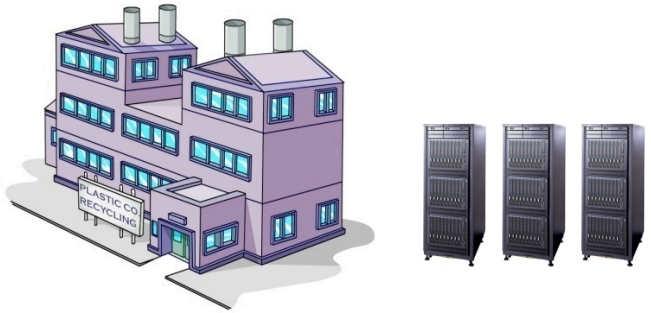

Host 4 Less

Snoopy's Startup

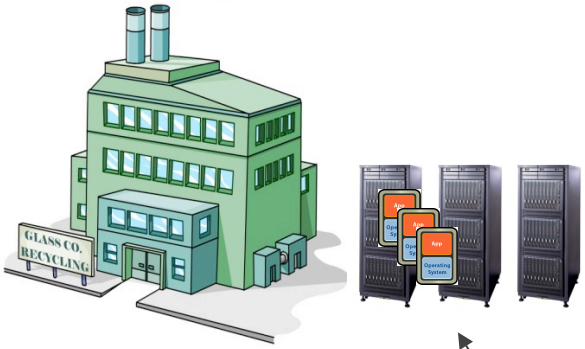

Host it R us.

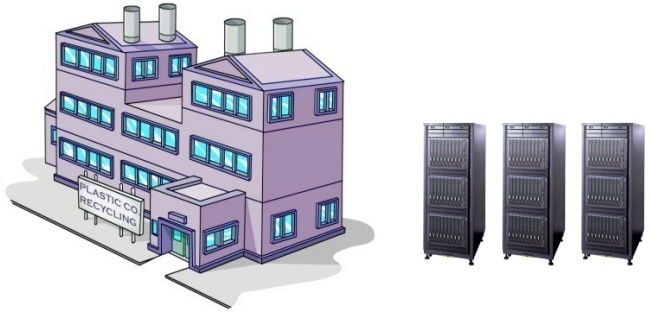

Host 4 Less

Dogspace

Snoopy's Startup

vmware

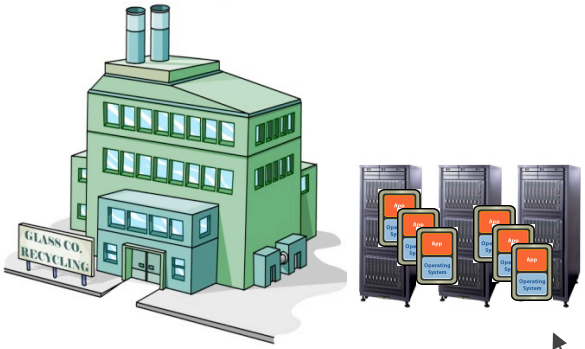

Host it R us.

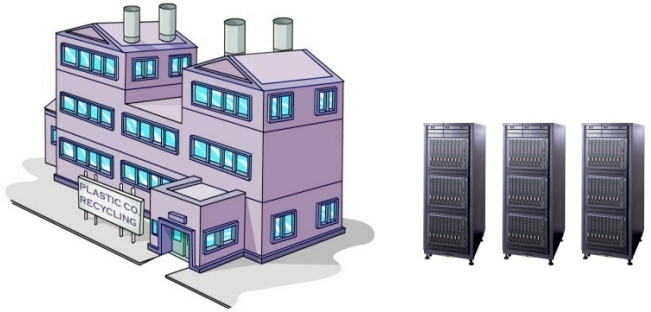

Host 4 Less

Dogspace

DogTube

Snoopy's Startup

vmware

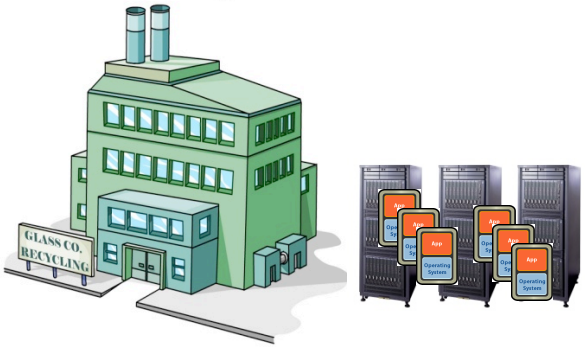

Host it R us.

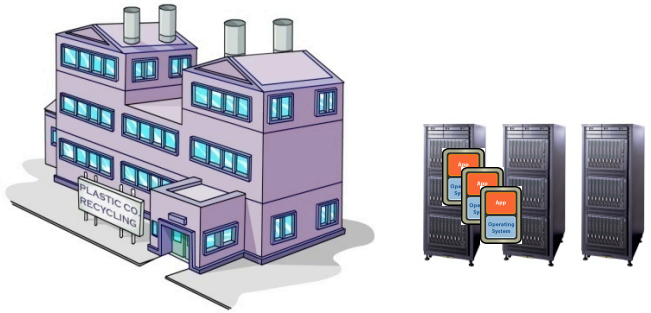

Host 4 Less

Dogspace

DogFlix

DogTube

Snoopy's Startup

vmware

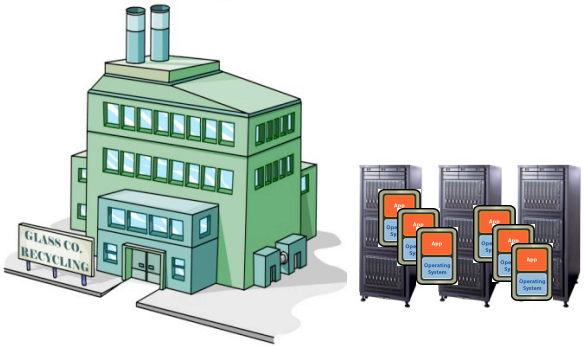

Host it R us.

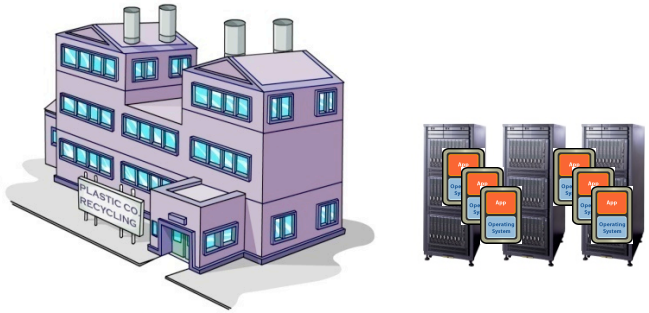

Host 4 Less

Dogspace

DogFlix

DogTube

dBay

Snoopy's Startup

vmware

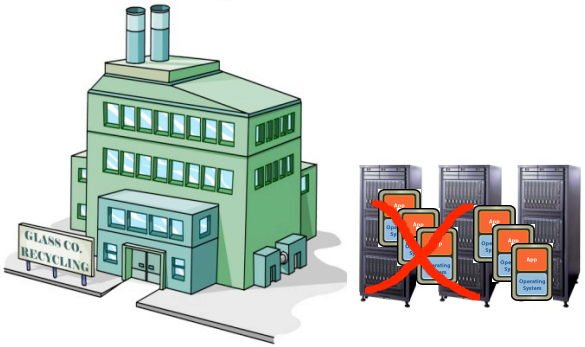

Host it R us.

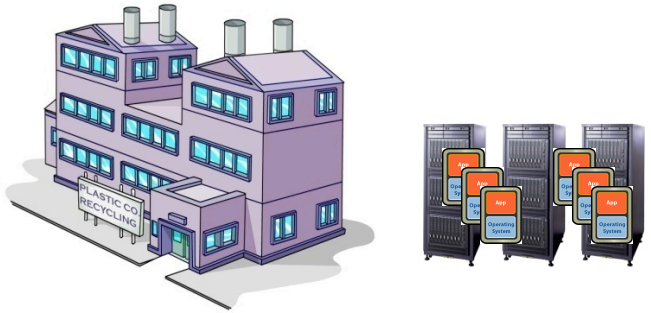

Host 4 Less

~~DogSpace~~

DogTube

DogFlix

dBay

Snoopy's Startup

vmware

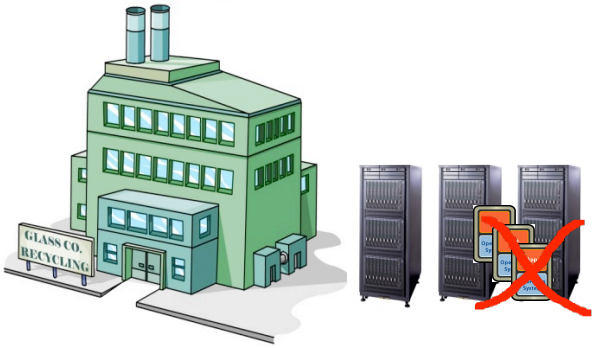

Host it R us.

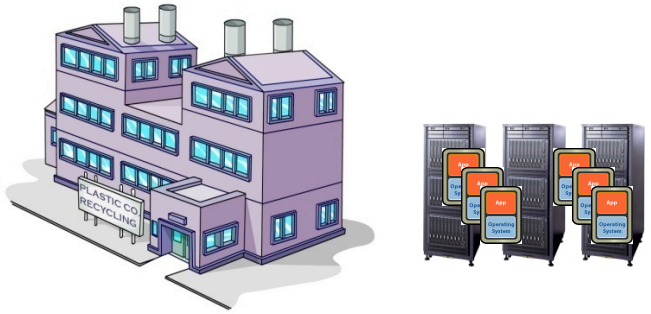

Host 4 Less

DogTube

DogFlix

dBay

Snoopy's Startup

vmware

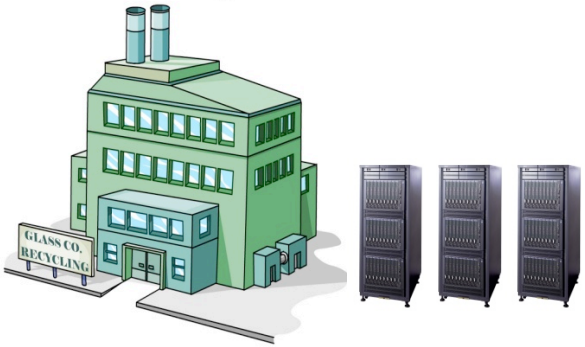

Host it R us.

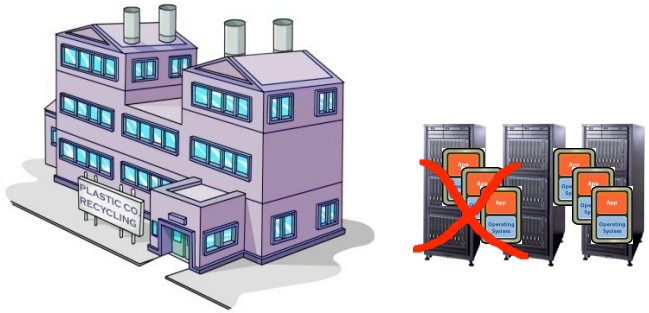

Host 4 Less

DogTube

dBay

Snoopy's Startup

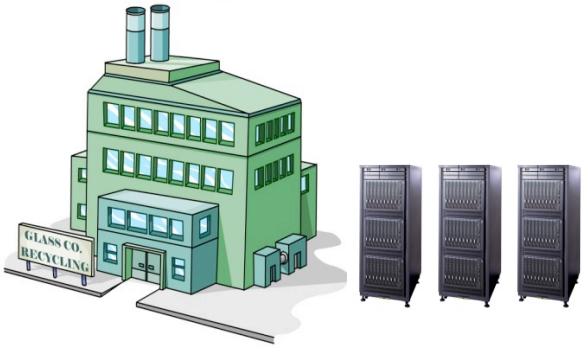

Host it R us.

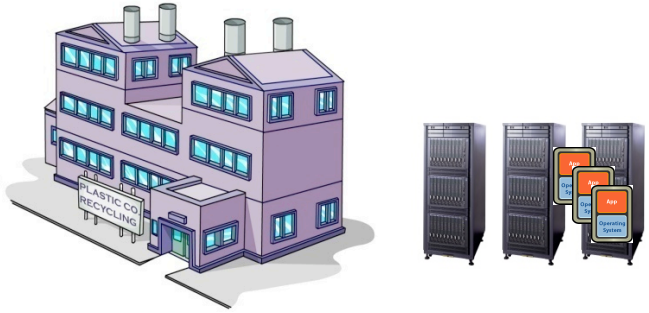

Host 4 Less

dBay 

Snoopy's Startup

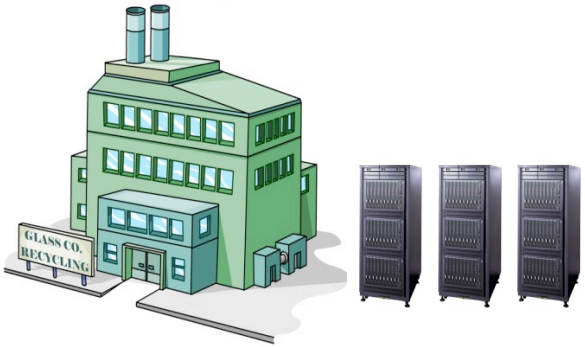

Host it R us.

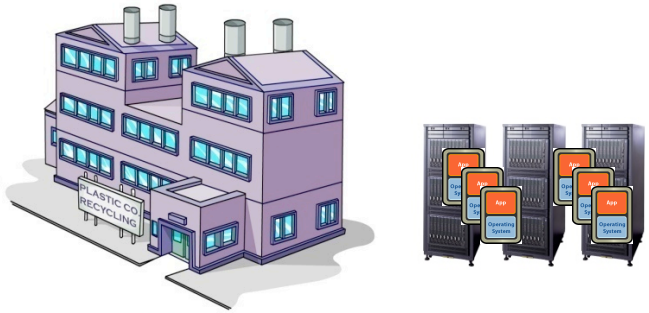

Host 4 Less

dBay

Snoopy's Startup

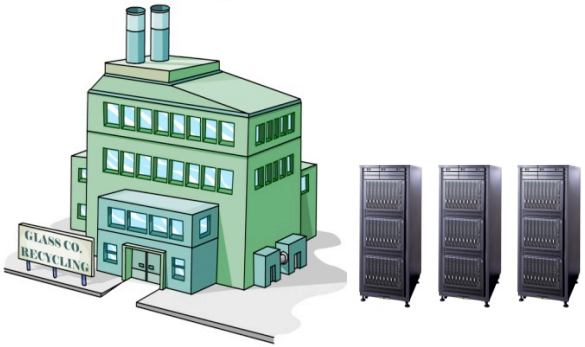

Host it R us.

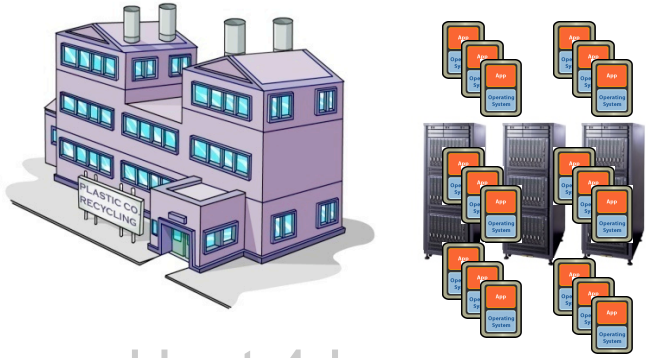

Host 4 Less

dBay

Snoopy's Startup

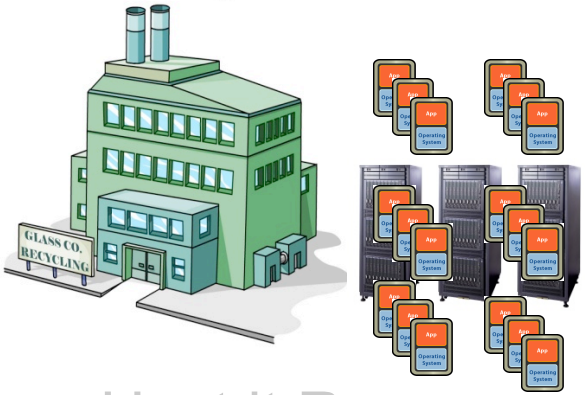

Host it R us.

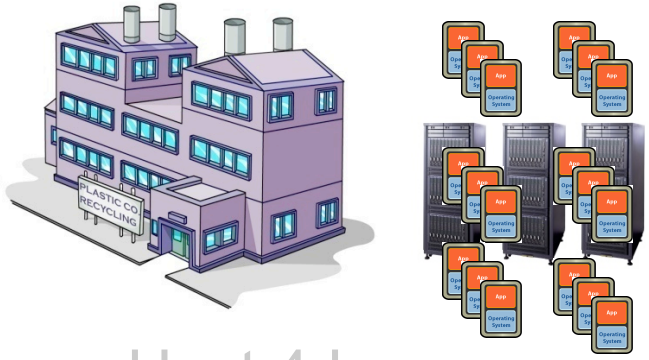

Host 4 Less

dBay

Snoopy's Startup

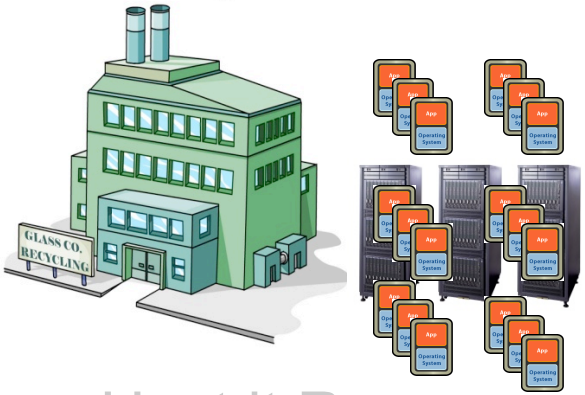

Host it R us.

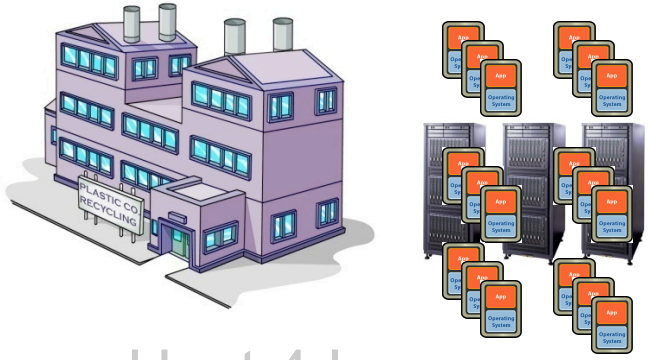

Host 4 Less

vmware®

What we want

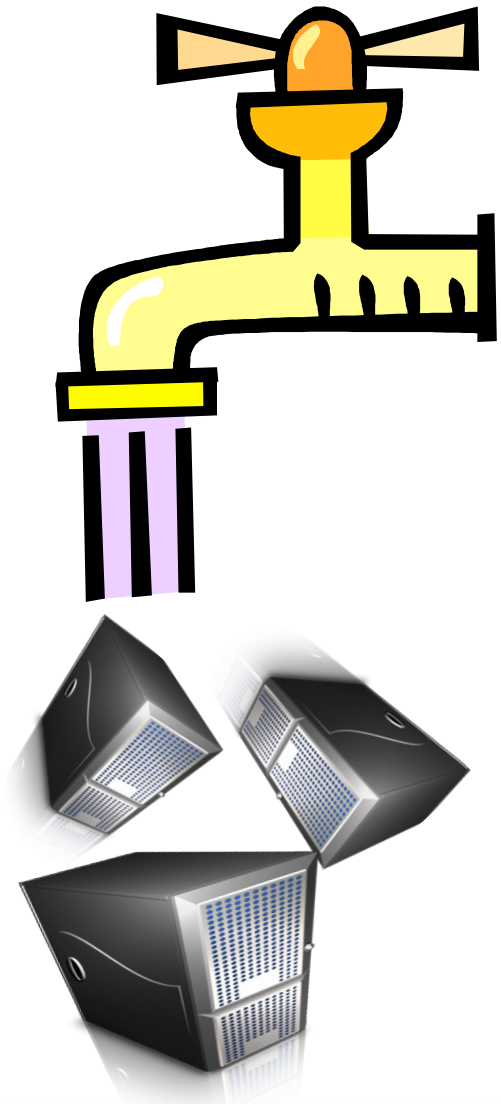

vmware®

This is really nothing new…

Original vision of Utility/grid computing:

"If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry."

John McCarthy, MIT Centennial in 1961

Virtualization converts computation into a fungible commodity

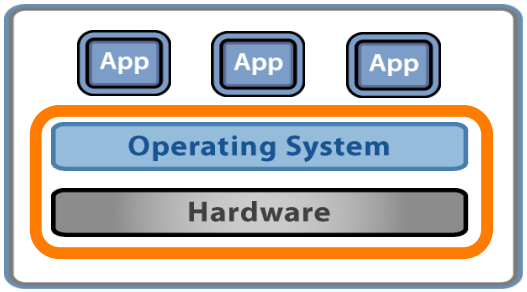

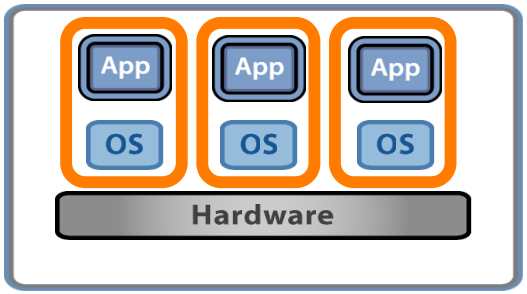

vmware

Two alternatives appear to be plausible

- **Small number of vertically integrated clouds:**
  - E.g., Amazon, MS, Google, IBM…
  - Highly differentiated in services and APIs
  - Each geographically distributed.
  - Each internally homogenous
  - Each with a broad community of ISVs
- **Marketplace of clouds: anyone can stand up their own cloud**
  - E.g., VMware, IBM, Eucalyptus, …
  - Many providers of consistent services, differentiated by price, service, geography, offerings…
  - Enables wide community to innovate in building clouds (e.g., researchers, HW vendors…)
  - Clouds may be built using very different techologies
  - Shared community of ISVs

vmware

Challenges to the marketplace model

- **Critical to define simple common abstractions independent implementation**
- **Need to support enterprise workloads, HPC workloads, … not just scale out web applications.**
- **Much harder to design:**
  - software that can be packaged.
  - software that can scale up and down.
  - a platform that enables third party innovation.

vmware®

Research Example: Implementation cloud management SW

- Instrumentation/metering of massive cloud
- Automation at scale
- Developing offerings for different use cases, e.g., university research
- Security: detection, not trusting service provider, ...
- Alternative implementations of vCloud API

vmware®

Example: Federation between clouds

- Exploiting long tail for data deduplication
- Catalog federated between clouds
- Encryption/security for data
- Trading floor/futures market for capacity
- Move VMs near users

vmware®

Example: Fungible computing

- **Virtualization lets us move anything anywhere, but your mileage may vary:**
  - Cost may be very different
  - Performance may be very different.
  - Cost to the SP may be very different
- **What metrics can be collected presented to user?**
- **How do we characterize application?**
- **How do we characterize physical capacity?**

Example: Enabling HPC

- Defining vDC optimized for HPC
- New scheduler for grid/batch tasks
- New scheduler for data intensive supercomputing: e.g., Hadoop
- Fork task across 100s of nodes: e.g. SnowFlock from UofToronto
- Enabling RDMA
- Special purpose OS?

vmware®

Example: Augmented desktop

- Follow me desktops
- Nested Security
- Remote apps vs. distributing apps vs. remote desktops – how to seamlessly integrate? (Unity)
- When your desktop is in the cloud, what changes?
- Disaster recovery
- Being a system admin for your mom.
- Moving from cloud provider to cloud provider.
- Remote 3d graphics over WAN – lots of challenges!

vmware

Example: Research in OS development

- **Scalable deployment for HPC.**
- **Communication protocol.**
- **Developing library OS that is re-usable.**
- **Scalability for massive multi-core.**
- **Migration to and from generic OS.**
  - Real time
  - Scalability
- **Control of TLB for managed code**
- **Code and file system sharing**
- **Examples: Managed Code (Libra IBM, Liquid VM BEA, Maxine VM SUN), HEC/HPC (UNM/Sandia, Cambridge) Games (Sony PS3), Denali from UofWashington**

vmware®

Example: The substrate

- **Whats the right cloud HW?**
  - Architecture can evolve without OS stifling innovation
  - All the end points are virtualized

- **How does the virtualization layer change?**
  - Implicitly managed storage hierarchy
  - Scale out storage
  - Networking services: FW, loadbalancer, VPN, …

vmware®

Concluding Remarks

- **Cloud computing is going to be transformative to our industry**
  - Transforms the Data Center
  - Transforms Clients
  - Transforms Management Techniques
  - Transforms Application Development
- **Lots of new research opportunities; nobody has tackled the tough problems.**
- **Vertically integrated Clouds will stifle research; Open clouds/ standards are critical for advancement and research**
- **Some parts of that research may be complementary to VMware's stack**
- **Open source vCloud API (www.vmware.com/go/vcloudapi)**