EC/CS 528: Cloud Computing

Introduction

Instructor: Alan Liu

BU Department of Electrical & Computer Engineering

BOSTON UNIVERSITY

About Alan?

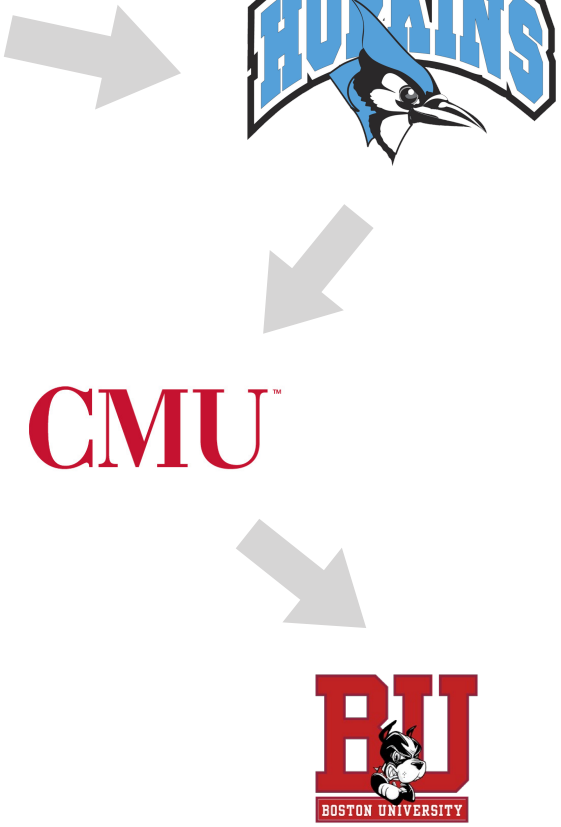

Industry Collaborators

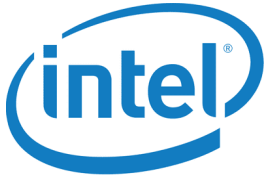

BOSTON UNIVERSITY

Introduction

- Instructor: Alan (Zaoxing) Liu
  - PhD in Computer Science from Johns Hopkins U.
  - Postdoc at Carnegie Mellon U.
  - Research areas: computer systems and networking
  - Current research
    - In-network and cloud computing
    - Data analytics systems
    - Software-defined networking
    - Network and cloud security
- Office hours: PHO 335, Tue 4pm or by appointment
- Teaching assistant: Zeying Zhu, Thu 7pm

BOSTON UNIVERSITY

What is Cloud Computing?

- What do you think?

- *"**Cloud computing** is an information technology (IT) paradigm that enables ubiquitous access to shared pools of configurable system resources and higher-level services that can be rapidly provisioned with minimal management effort, often over the Internet. Cloud computing relies on sharing of resources to achieve coherence and economies of scale, similar to a public utility."*
  https://en.wikipedia.org/wiki/Cloud_computing

- *"**Cloud computing** is the delivery of computing services – servers, storage, databases, networking, software, analytics and more – over the Internet ("the cloud"). Companies offering these computing services are called cloud providers and typically charge for cloud computing services based on usage, similar to how you're billed for gas or electricity at home."*
  https://azure.microsoft.com/en-gb/overview/what-is-cloud-computing/

BOSTON UNIVERSITY

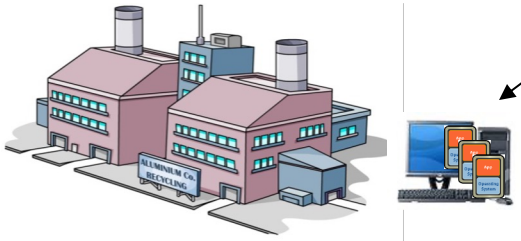

IT dept.

Joe’s Widget Co.

BOSTON UNIVERSITY

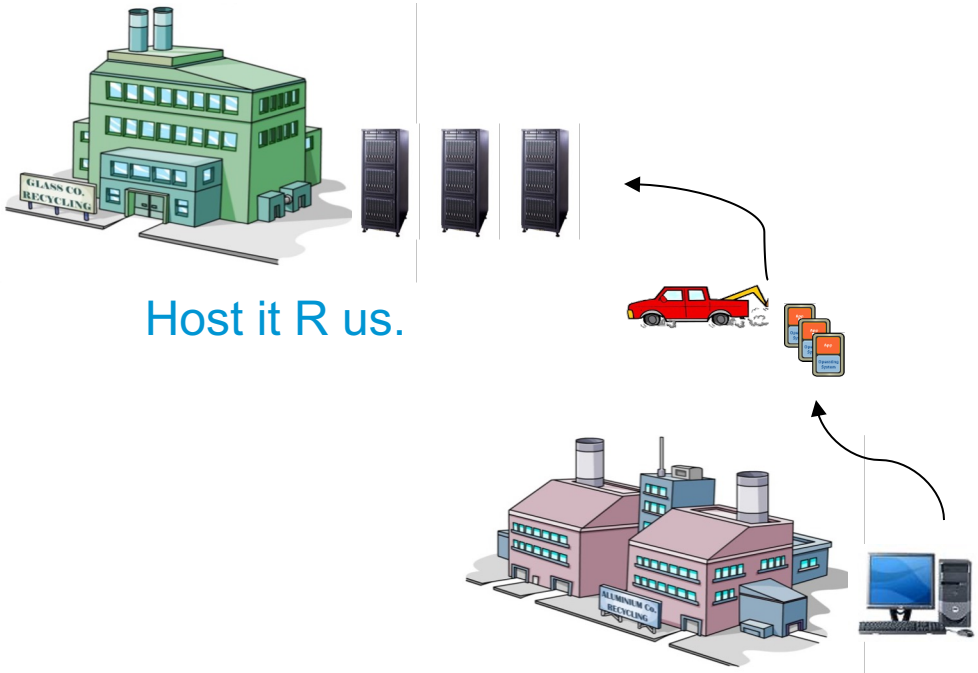
GLASS CO.
RECYCLING
Host it R us.
ALUMINUM CO.
RECYCLING
Joe's Widget Co.

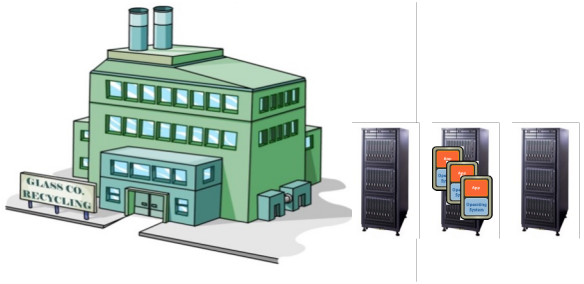

Host it R us.

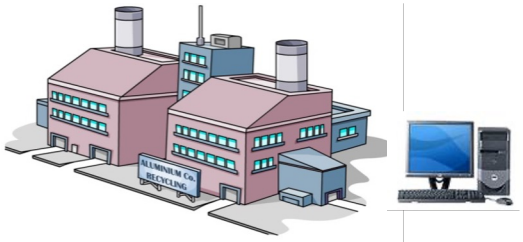

Joe’s Widget Co.

BOSTON UNIVERSITY

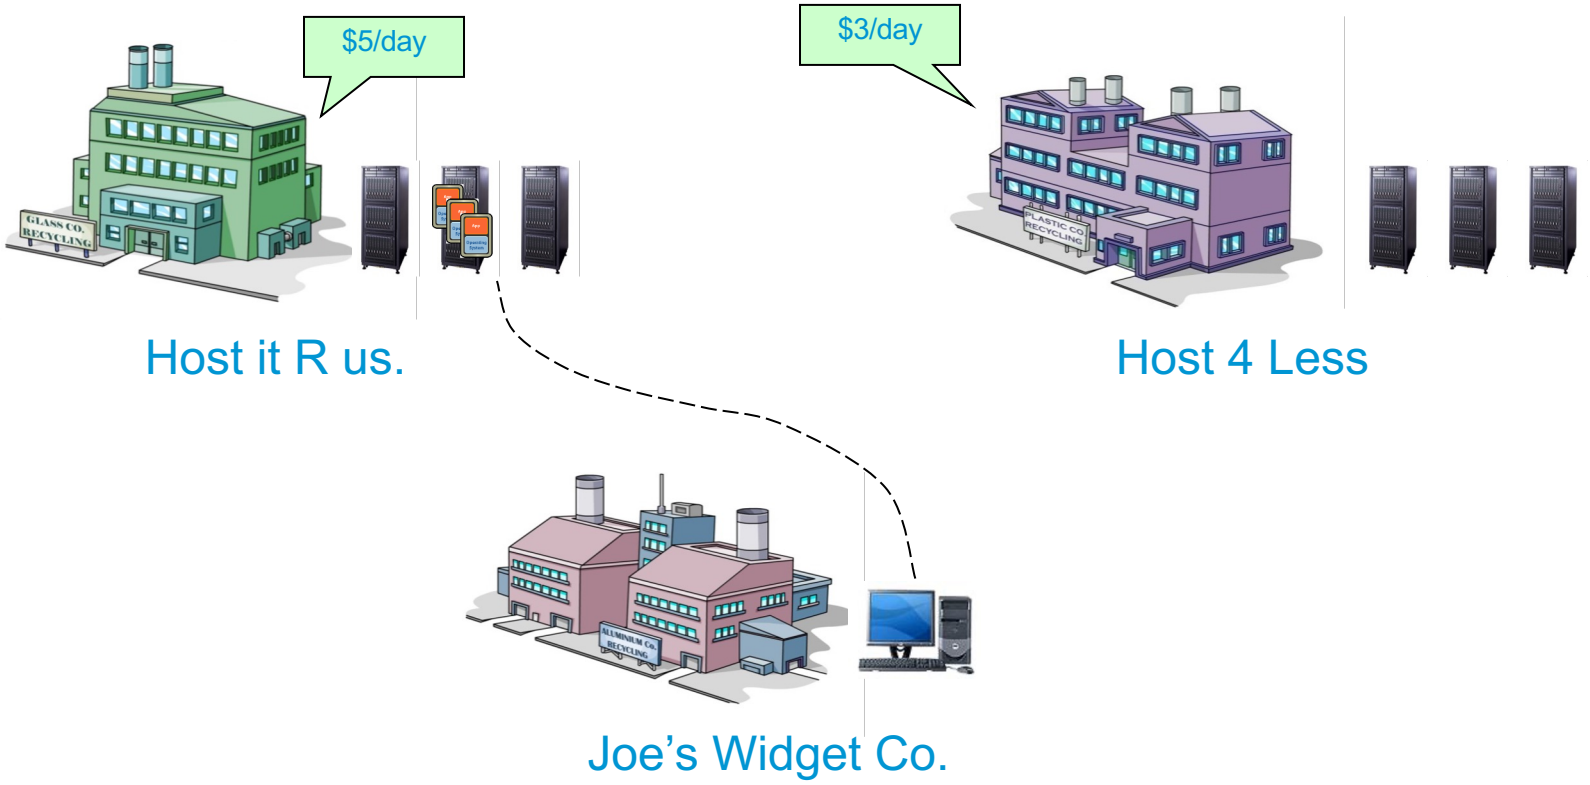

BOSTON UNIVERSITY

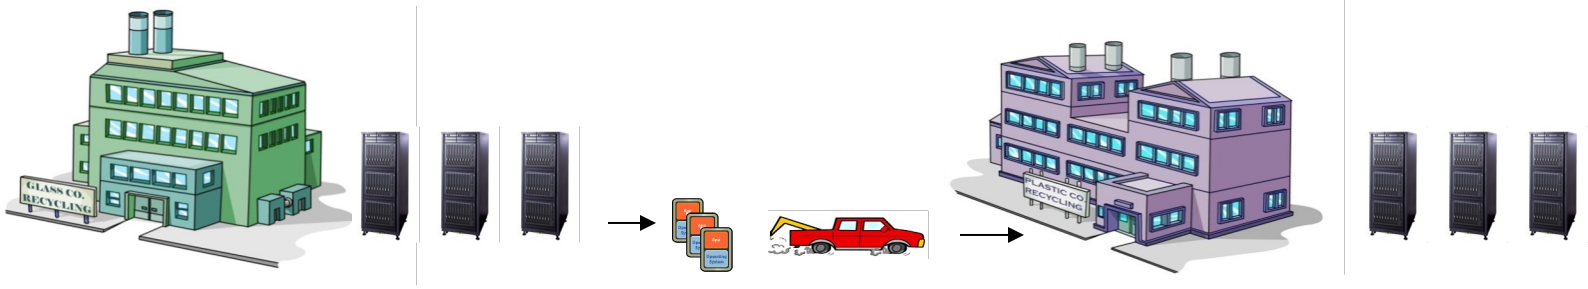

Host it R us.

Host 4 Less

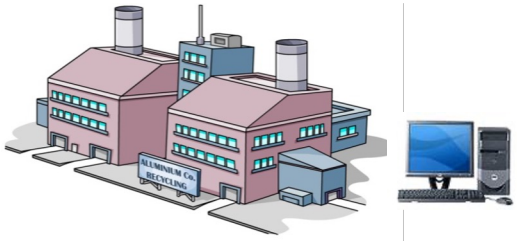

Joe’s Widget Co.

BOSTON UNIVERSITY

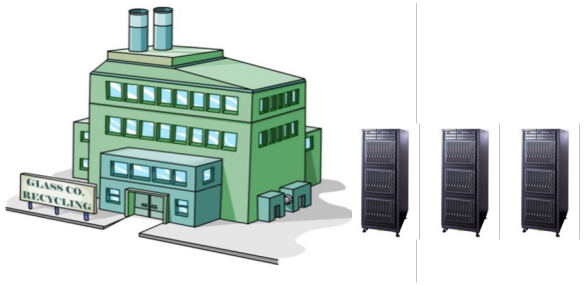

Host it R us.

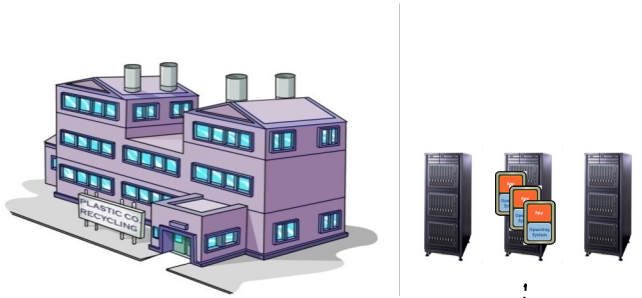

Host 4 Less

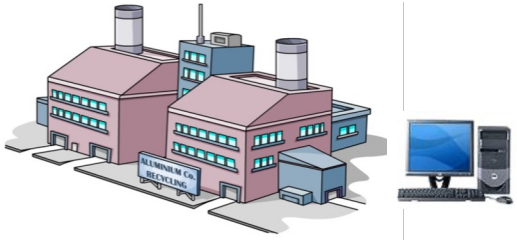

Joe's Widget Co.

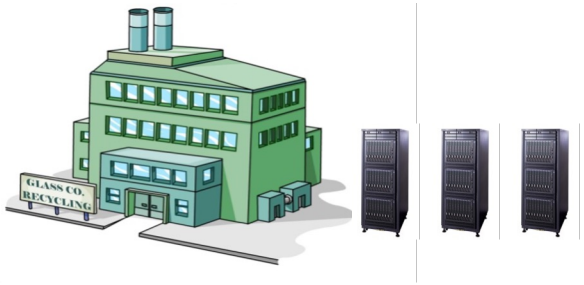

Host it R us.

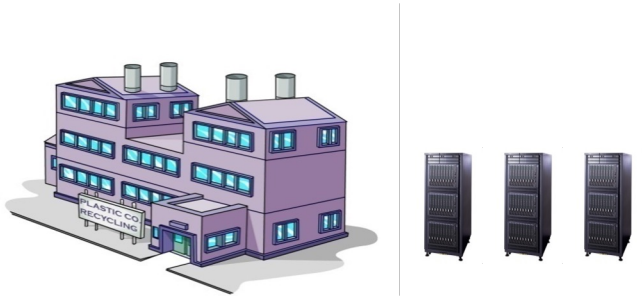

Host 4 Less

Snoopy's Startup

BOSTON UNIVERSITY

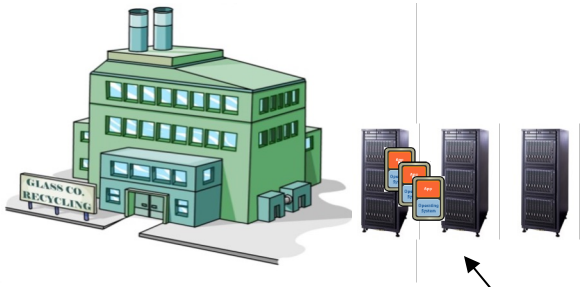

Host it R us.

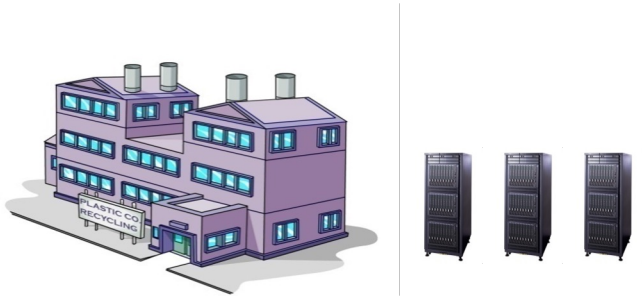

Host 4 Less

Dogspace

Snoopy's Startup

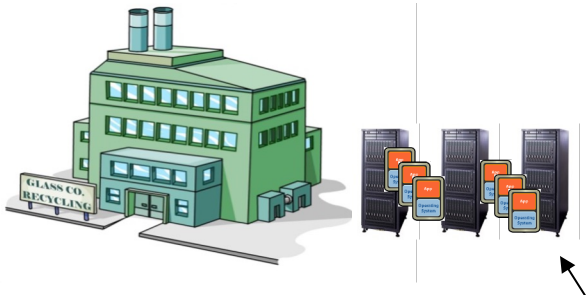

Host it R us.

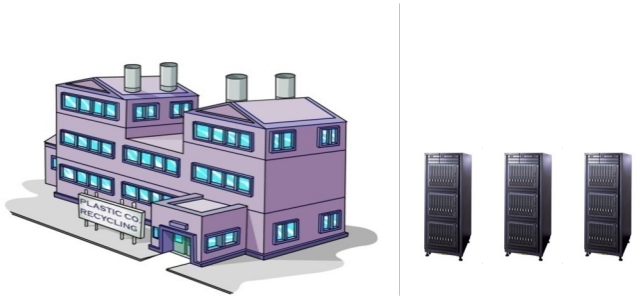

Host 4 Less

Dogspace

DogTube

Snoopy's Startup

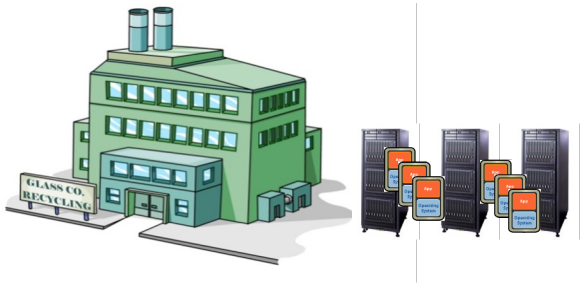

Host it R us.

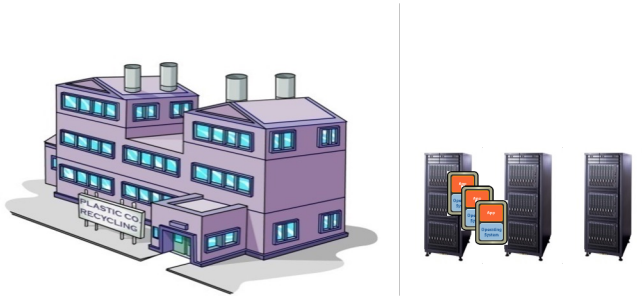

Host 4 Less

Dogspace

DogFlix

DogTube

Snoopy's Startup

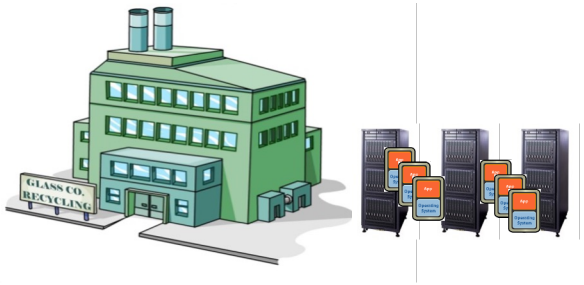

Host it R us.

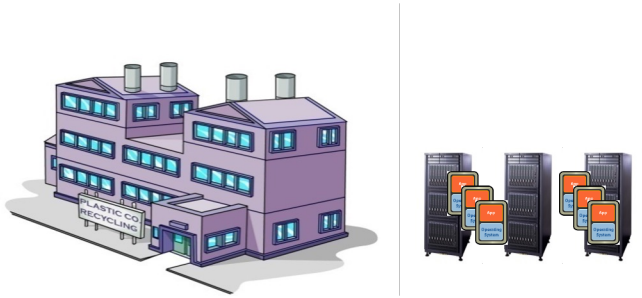

Host 4 Less

Dogspace

DogFlix

DogTube

dBay

Snoopy's Startup

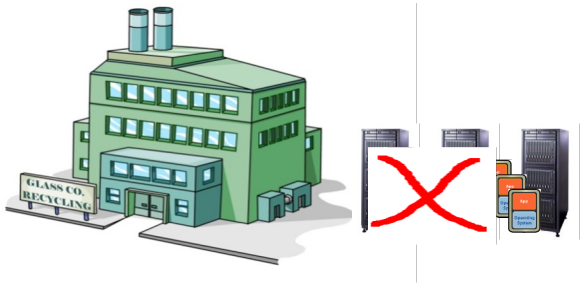

Host it R us.

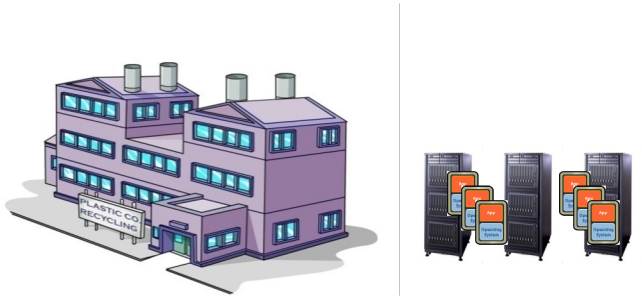

Host 4 Less

DogFlix

Snoopy's Startup

DogTube

dBay

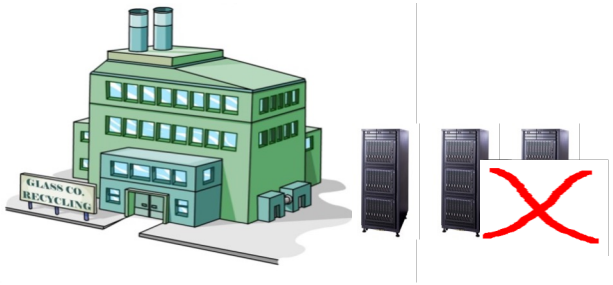

Host it R us.

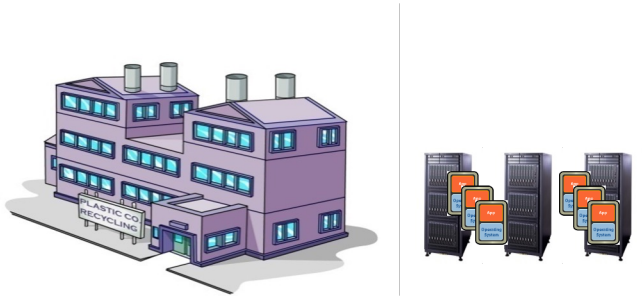

Host 4 Less

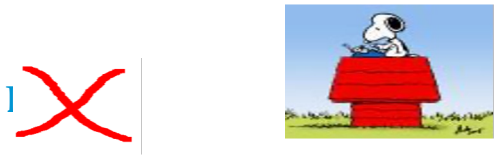

Snoopy's Startup

DogTube

dBay

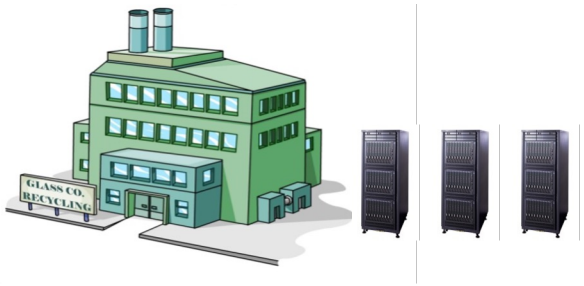

Host it R us.

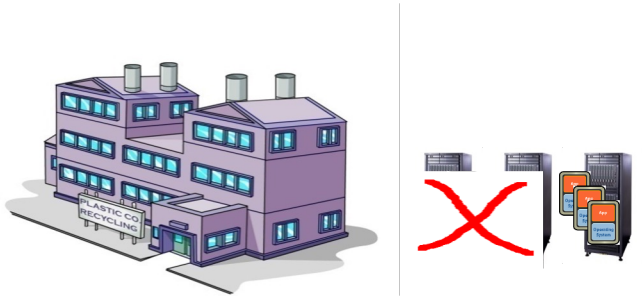

Host 4 Less

Snoopy's Startup

dBay

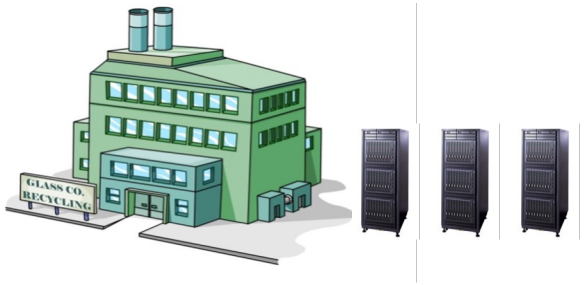

Host it R us.

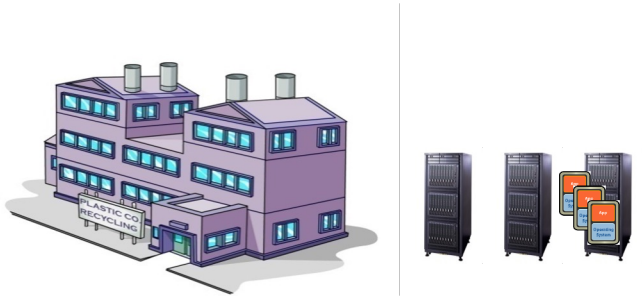

Host 4 Less

Snoopy's Startup

dBay

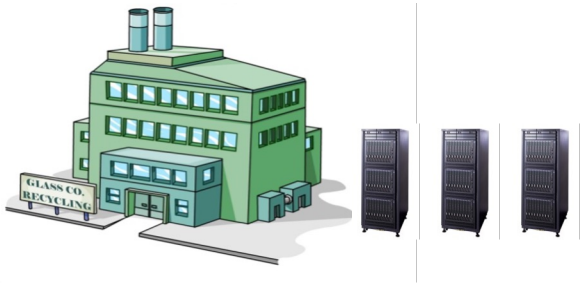

Host it R us.

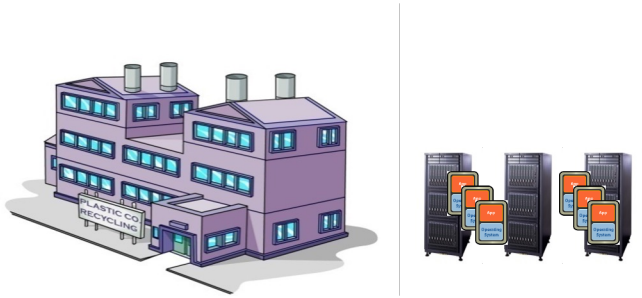

Host 4 Less

Snoopy's Startup

dBay

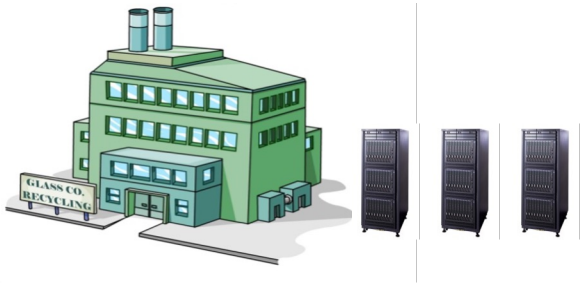

Host it R us.

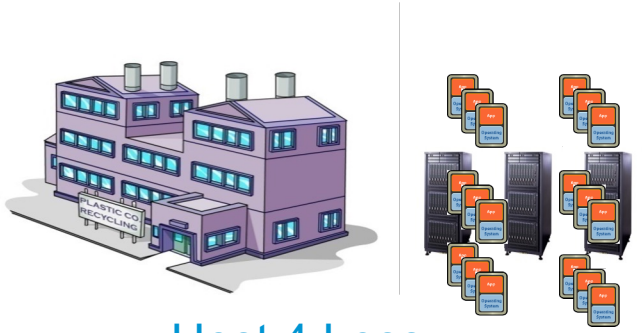

Host 4 Less

Snoopy's Startup

dBay

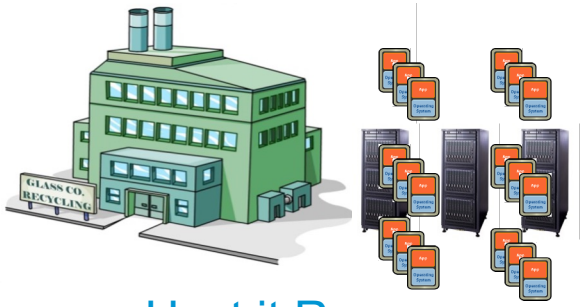

Host it R us.

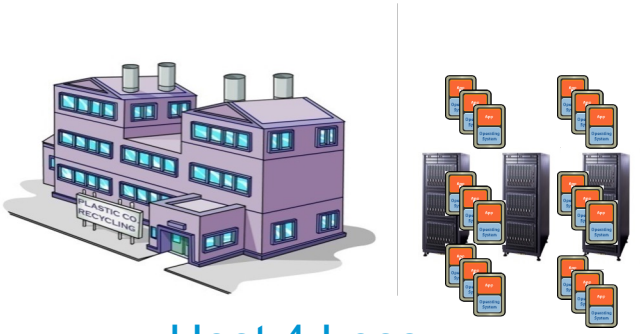

Host 4 Less

Snoopy's Startup

dBay 

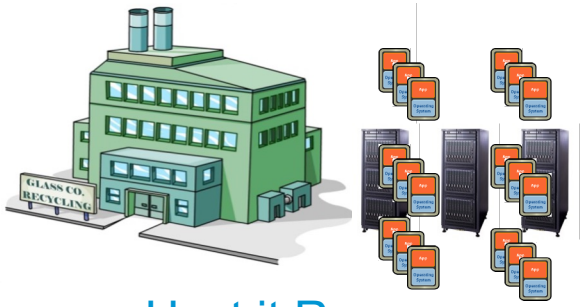

Host it R us.

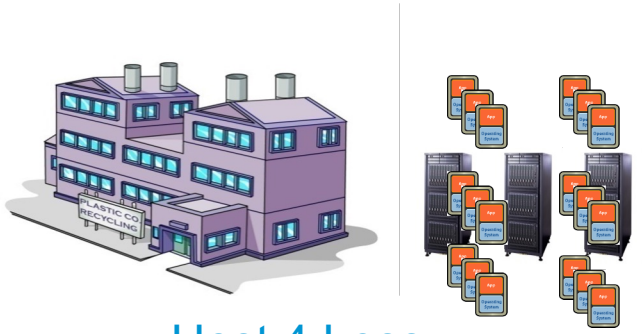

Host 4 Less

What we want?

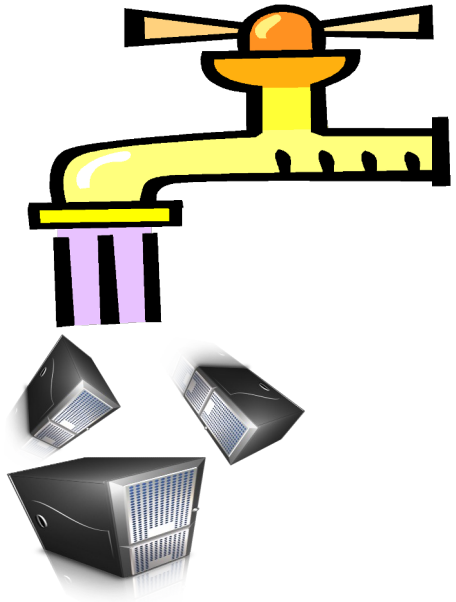

Cloud Computing Models, Resources, Attributes

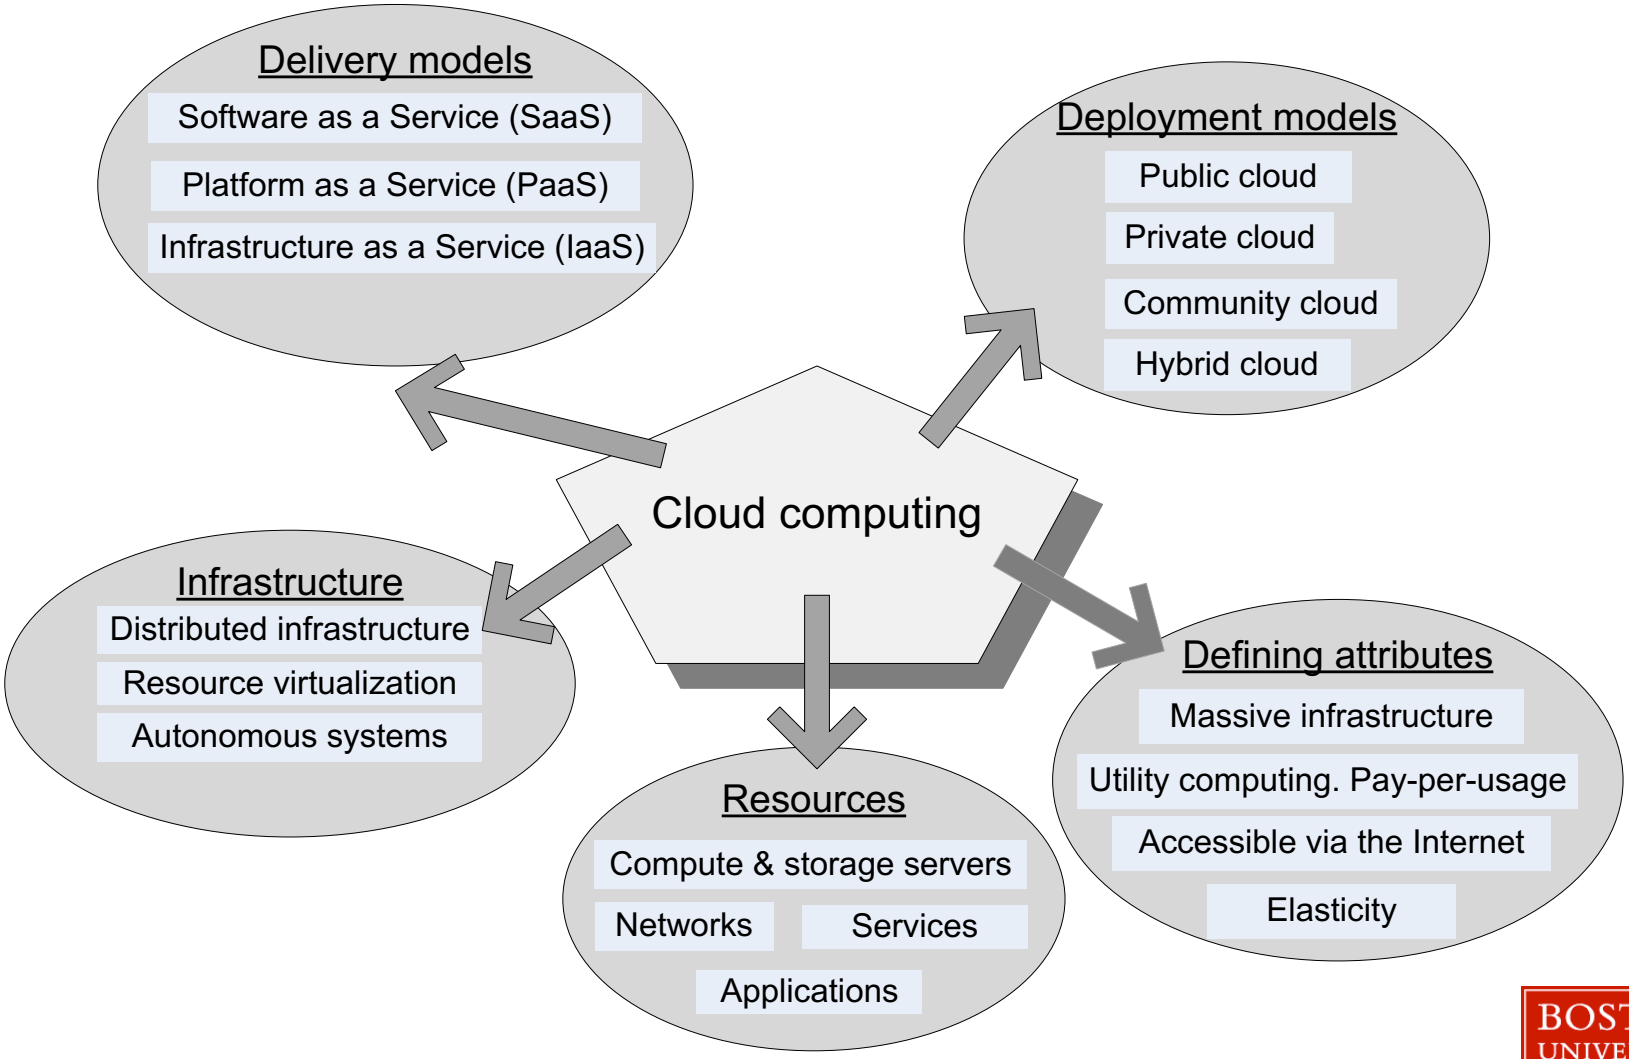

BOSTON UNIVERSITY

Types of Clouds

1. **Public Cloud** - the infrastructure is made available to the general public or a large industry group and is owned by the organization selling cloud services.

2. **Private Cloud** – the infrastructure is operated solely for an organization.

3. **Hybrid Cloud** - composition of two or more Clouds (public, private, or community) as unique entities but bound by a standardized technology that enables data and application portability.

4. **Other types: e.g., Community/Federated Cloud** - the infrastructure is shared by several organizations and supports a community that has shared concerns.

BOSTON UNIVERSITY

Course Structure

- Series of Agile lectures: tentatively by Michael Daitzman
- Lectures on Cloud Computing
- Invited talks
- Projects mentored by the industry

You can get to course webpage. This will be kept up to date, has all deadlines, demos…

Piazza is the source of truth.

BOSTON UNIVERSITY

Course Logistics

- Course webpage
  - https://zaoxing.github.io/teaching/2022-fall-cloud
  - Paper readings.
  - Updates on the schedule (look for Piazza signup code!)

- Capstone project
  - A semester-long project with 3~5 students.
  - Mentored by external researchers and engineers!
    - Intel, Microsoft, OpenInfra Foundation, Snowflake, PingCAP, Bell Labs, BCH, etc.
  - Access public cloud resources
    - Mass Open Cloud, https://massopen.cloud/
    - AWS, Azure, GCP

BOSTON UNIVERSITY

Course Logistics

- Project matching
  - We will send out a user skill survey this week.

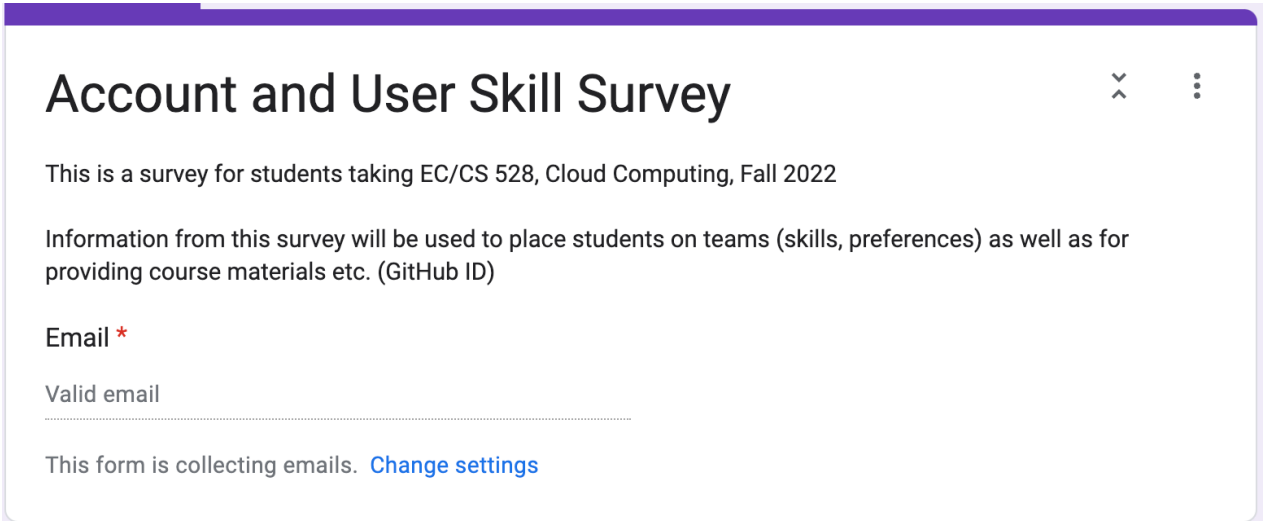

- Paper readings
  - 1~2 readings per week.
  - No need to write reviews, but you need to read them,
  - May have random quizzes via Piazza.

BOSTON UNIVERSITY

Sprint project presentations

- Every two weeks, all groups present/demo their projects.

- By 9AM the morning before demos, you will submit a 12 minute video describing your progress, and a quiz for your fellow students:
  - remember you are teaching us about the project and what you have learned
  - minimize acronyms
  - try to get everyone to talk
  - try to make sure you review it with the mentor

- We will randomly select a set of the presentations, ask the team to come to the front, and we and the other students will ask you questions about it.

- We will provide feedback on the projects that have presented by the next day.

BOSTON UNIVERSITY

How you are evaluated

- Project: 70 %
  - project description (proposal → report) 10% (template & example)
  - bi-weekly project demo and verbal status in aggregate 30%
  - final project result and report (evolved description) 25%
  - project quizzes 5%
  - project scores are scaled based on team & mentor evaluation
- Paper quizzes & Piazza discussions: 30%
  - Paper quizzes will be run before or in the first few minutes of class
    - PhD students all assigned papers.
    - Masters and undergraduates will have reduced workloads.

BOSTON UNIVERSITY

But, Why Should You Stay in This Class, and This Field?

BOSTON UNIVERSITY

We will discuss

- This is the future of all computing
- It has profound impact on:
  - how hardware is developed
  - how software is developed
  - the computer industry
  - society as a whole

BOSTON UNIVERSITY

Course contents: a variety of topics

- Overview of Cloud Computing
- Virtualization
- Agile development
- Distributed systems/batch systems
- Datacenter architectures
- Storage systems
- Resource management
- Big data processing
- Datacenter networking
- Resource disaggregation
- Other research topics

BOSTON UNIVERSITY

So, Why is Cloud Computing Cool?

- Shared Resources and Resource Management
  - Cloud uses a shared pool of resources
  - Uses Internet to offer **scalable** and **elastic** services.
  - The term "**elastic computing**" refers to the ability of dynamically and on-demand acquiring computing resources and supporting a variable workload.
  - Resources are metered and users are charged accordingly.
  - It is more cost-effective due to resource-multiplexing. Lower costs for the cloud service provider are past to the cloud users.

Cloud Infra might be running under the water

BOSTON UNIVERSITY

Cloud Computing Benefits?

- Resources, such as CPU cycles, storage, network bandwidth, are shared.
- When multiple applications share a system, their peak demands for resources are not synchronized thus, **multiplexing** leads to a higher resource utilization.
- Resources can be **aggregated** to support data-intensive applications.
- Data sharing facilitates **collaborative** activities. Many applications require multiple types of analysis of shared data sets and multiple decisions carried out by groups scattered around the globe.

Cloud Computing Benefits (cont'd)

- Eliminates the **initial investment costs** for a private computing infrastructure and the maintenance and operation costs.
- **Cost reduction:** concentration of resources creates the opportunity to pay as you go for computing.
- **Elasticity:** the ability to accommodate workloads with very large peak-to-average ratios.
- **User convenience:** virtualization allows users to operate in familiar environments rather than in idiosyncratic ones.

BOSTON UNIVERSITY

Emulating the real world

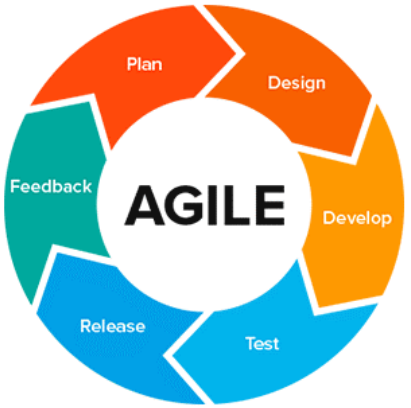

- Agile software development team.
- Learning by doing.
  - How to use cloud resources.
  - Programming and running real systems.
- Many of the projects relate to
  - state-of-the-art systems running in the cloud.
  - latest systems research.
- There are all kinds of internship opportunities and ways to engage if you are doing well!

BOSTON UNIVERSITY

Some Suggestions

- Working as a team:
  - Some of your team members will flake out
  - Some of your mentors will flake out
  - Some of your mentors will have unrealistic expectations
- Be in touch with your mentors.
- Let the instructors know early if there are problems
- Please take this seriously:
  - You are representing BU to mentors!
  - You will be doing a final public demo/presentation with industry invited.

BOSTON UNIVERSITY

How to Read

You May Think You Already Know How To READ, But…

BOSTON UNIVERSITY

We Spend a Lot of Time Reading

- Reading for classes
- Reviewing conference submissions
- Giving colleagues feedback
- Keeping up with your field
- Staying broadly educated
- Transitioning into a new areas
- Learning how to write better papers

It is worthwhile to learn to read *effectively*

BOSTON UNIVERSITY

Keshav's Three-Pass Approach: Step 1

- A ten-minute scan to get the general idea
  - Title, abstract, and introduction
  - Section and subsection titles
  - Conclusion and bibliography
- What to learn: the five C's
  - Category: What type of paper is it?
  - Context: What body of work does it relate to?
  - Correctness: Do the assumptions seem valid?
  - Contributions: What are the main research contributions?
  - Clarity: Is the paper well-written?
- Decide whether to read further…

BOSTON UNIVERSITY

Keshav's Three-Pass Approach: Step 2

- A more careful, one-hour reading
  - Read with greater care, but ignore details like proofs
  - Figures, diagrams, and illustrations
  - Mark relevant references for later reading
- Grasp the content of the paper
  - Be able to summarize the main idea
  - Identify whether you can (or should) fully understand
- Decide whether to
  - Abandon reading in greater depth
  - Read background material before proceeding further
  - Persevere and continue for a third pass

Keshav's Three-Pass Approach: Step 3

- Several-hour virtual re-implementation of the work
  - Making the same assumptions, recreate the work
  - Identify the paper's innovations and its failings
  - Identify and challenge every assumption
  - Think how you would present the ideas yourself
  - Jot down ideas for future work
- When should you read this carefully?
  - Reviewing for a conference or journal
  - Giving colleagues feedback on a paper
  - Understanding a paper closely related to your research
  - Deeply understanding a classic paper in the field

BOSTON UNIVERSITY

Other Tips for Reading Papers

- Read at the right level for what you need
  - "Work smarter, not harder"
- Read at the right time of day
  - When you are fresh, not sleepy
- Read in the right place
  - Where you are not distracted, and have enough time
- Read actively
  - With a purpose (what is your goal?)
  - With a pen or computer to take notes
- Read critically
  - Think, question, challenge, critique, …

BOSTON UNIVERSITY

Near-term deadlines

- Reading for next week: The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines (Only Chapters 1 to 3).

- If dropping the course: before signing up projects.

- Skills/project survey by weekend, will be posted on Piazza - BE HONEST!
- Project selection: before Wed 9/14 at 9AM - i.e., preferred drop date
  - You will select your top X projects, we will assign you to a team: pick projects you have background.
  - We will let you know teams by Wed 9/14 evening
  - Any students joining after 9/14 skill survey only
  - Meet with you mentor… ASAP
- Project plan/proposal 1 week later: Wednesday 9/21
- First demo will be following week (you won't know if you are doing a demo)

BOSTON UNIVERSITY

Self-introduction

Q&A

BOSTON UNIVERSITY