
CT4I4

Distributed Systems & Co-Operative Computing

Name: Andrew Hayes
Student ID: 21321503
E-mail: a.hayes18@universityofgalway.ie

2025-03-12

Contents

1	Introduction	1
1.1	Client-Server Architectures	1
1.1.1	Two-Tier Architectures	1
1.2	Three-Tier Architecture	1
1.3	Network Programming Paradigms	1
2	Java RMI	2
2.1	Steps to Creating an RMI Application	2
2.2	Example Java RMI Program	3
3	Enterprise Java Beans	7
3.1	Distributed System Scenario	7
3.2	EJB	7
3.2.1	Entity EJBs	9
4	NodeJS	9
4.1	MEAN	10
5	Proxmox Virtualisation Environment	10
5.1	KVM	11
5.2	QEMU	11
5.3	LXC	11
5.4	Ceph	11
5.4.1	Ceph Network	12
5.4.2	Ceph OSDs	13
5.5	VM Installation	13
5.5.1	Hard Drives	13
5.5.2	Memory	13
5.5.3	VM Backups	13
5.5.4	VM Migration	13
5.5.5	VM Cloning	13
5.5.6	VM Imports	14
5.5.7	User Authentication	14
5.6	Proxmox Cluster	14
5.7	High Availability	14
5.7.1	HA Groups	15
5.8	Performance Benchmarking	15
6	Cloud Computing	15

1 Introduction

1.1 Client-Server Architectures

1.1.1 Two-Tier Architectures

A **two-tier client-server architecture** is a client-server architecture wherein a client talks directly to a server, with no intervening server. It is typically used in small environments ($\lesssim 50$ users).

A common development error is to prototype an application in a small, two-tier environment, and then scale up by simply adding more users to the server: this approach will usually result in an ineffective system, as the server becomes overwhelmed. To properly scale to hundreds or thousands of users, it is usually necessary to move to a three-tier architecture.



Figure 1: Client & server using TCP/IP protocols to communicate. Information can flow in either or both directions. The client & server can interact with a transport layer protocols.

1.2 Three-Tier Architecture

A **three-tier client-server architecture** introduces a server or **agent** (or **load-balancer**) between the client & the server. The agent has many roles:

- Translation services: such as adapting a legacy application on a mainframe to a client-server environment.
- Metering services: such as acting as a transaction monitor to limit the number of simultaneous requests to a given server.
- Intelligent agent services: as in mapping a request to a number of different servers, collating the results, and returning a single response to the client.

1.3 Network Programming Paradigms

Practically all network programming is based on a client-server model; the only real difference in paradigms is the **level** at which the programmer operates. The sockets API provides direct access to the available transport layer protocols. RPC is a higher-level abstraction that hides some of the lower-level complexities. Other approaches are also possible:

- Sockets are probably the best-known and most widely-used paradigm. However, problems of data incompatibility across platforms can arise.
- RPC libraries aim to solve some of the basic problems with sockets and provide a level of transport independence.
- Neither approach works very well with modern applications (Java RMI and other modern technologies, e.g., web services are better).

2 Java RMI

Remote Method Invocation (RMI) is a Java-based mechanism for distributed object computing. RMI enables the distribution of work to other Java objects residing in other processes or on other machines. The objects in one Java Virtual Machine (JVM) are allowed to seamlessly invoke methods on objects in a remote JVM. To call a method of a remote object, we must first get a reference to that object, which can be obtained from the registry name facility or by receiving the reference as an argument or return value of a method call. Clients can call a remote object in a server that itself is a client of another server. Parameters of method calls are passed as serialised objects:

- types are not truncated, and therefore, object-oriented polymorphism is supported;
- parameters are passed by value (deep copy) and therefore object behaviour can be passed.

The Java Object Model is still supported with distributed (remote) objects. A reference to a remote object can be passed to or returned from local & remote objects. Remote object references are passed by reference: therefore, the whole object is not always downloaded. Objects that implement the `Remote` interface are passed as a remote reference, while other objects are passed by value (using object serialisation).




Figure 2: Java RMI Architecture

The client obtains a reference for a remote object by calling `Naming.lookup(//URL/registered_name)` which is a method which returns a reference to another remote object. Methods of the remote object may then be called by the client. This call is actually to the **stub** which represents the remote object. The stub packages the arguments (**marshalling**) into a data stream (to be sent across the network). On the implementation side, the skeleton unmarshals the argument, calls the method, marshals the return value, and sends it back. The stub unmarshals the return value and returns it to the caller. The RMI layer sits on top of the JVM and this allows it to use Java Garbage Collection of remote objects, Java Security (a security manager may be set for the server, now deprecated), and Java class loading.

2.1 Steps to Creating an RMI Application

1. Define the interfaces to your remote objects.
2. Implement the remote object classes.
3. Write the main client & server programs.

4. Create the stub & skeleton classes by running the *rmic* compiler on the remote implementation classes. (No longer needed in later Java versions).
5. Start the *rmiregistry* (if not already started).
6. Start the server application.
7. Start the client (which contains some initial object references).
8. The client application/applet may then call object methods in the remote (server) program.

2.2 Example Java RMI Program

```

1 // Remote Object has a single method that is passed
2 // the name of a country and returns the capital city.
3 import java.rmi.*;
4
5 public interface CityServer extends Remote
6 {
7     String getCapital(String Country) throws
8         RemoteException;
9 }

```

Listing 1: Example Java RMI Program

```

1 import java.rmi.*;
2 import java.rmi.server.*;
3
4 public class CityServerImpl
5     extends UnicastRemoteObject
6     implements CityServer
7 {
8     // constructor is required in RMI
9     CityServerImpl() throws RemoteException
10    {
11        super(); // call the parent constructor
12    }
13
14    // Remote method we are implementing!
15    public String getCapital(String country) throws
16        RemoteException
17    {
18        System.out.println("Sending return string now - country requested: " + country);
19        if (country.toLowerCase().compareTo("usa") == 0)
20            return "Washington";
21        else if (country.toLowerCase().compareTo("ireland") == 0)
22            return "Dublin";
23        else if (country.toLowerCase().compareTo("france") == 0)
24            return "Paris";
25        return "Don't know that one!";
26    }
27
28    // main is required because the server is standalone
29    public static void main(String args[])

```

```

30     {
31         try
32         {
33             // First reset our Security manager
34             System.setSecurityManager(new RMISecurityManager());
35             System.out.println("Security manager set");
36
37             // Create an instance of the local object
38             CityServerImpl cityServer = new CityServerImpl();
39             System.out.println("Instance of City Server created");
40
41             // Put the server object into the Registry
42             Naming.rebind("Capitals", cityServer);
43             System.out.println("Name rebind completed");
44             System.out.println("Server ready for requests!");
45         } catch(Exception exc)
46         {
47             System.out.println("Error in main - " + exc.toString());
48         }
49     }
50 }

```

Listing 2: Example Server Implementation

```

1  public class CityClient
2  {
3      public static void main (String args[])
4      {
5          CityServer cities = (CityServer) Naming.lookup("//localhost/Capitals");
6          try {
7              String capital = cities.getCapital("USA");
8              System.out.println(capital);
9          } catch (Exception e) {}
10     }
11 }

```

Listing 3: Example Client Implementation

No distributed system can mask communication failures: method semantics should include failure possibilities. Every RMI remote method must declare the exception `RemoteException` in its **throw** clause. This exception is thrown when method invocation or return fails. The Java compiler requires the failures to be handled.

When implementing a remote object, the implementation class usually extends the RMI class `UnicastRemoteObject`: this indicates that the implementation class is used to create a single (non-replicated) remote object that uses RMI's default sockets-based transport for communication. If you choose to extend a remote object from a non-remote class, you need to explicitly export the remote object by calling the method `UnicastRemoteObject.exportObject()`.

The main method of the service first needs to create & install a **security manager**, either the `RMISecurityManager` or one that you have defined yourself. A security manager needs to be running so that it can guarantee that the classes loaded do not perform “sensitive” operations. If no security manager is specified, no class loading for RMI classes is allowed, local or otherwise.

TO make classes available via a web server (or your classpath), copy them into your public HTML directory. Alternatively, you could have compiled your files directly into your public HTML directory:

```

1 javac -d ~/project_dir/public_html City*.java
2 rmic -d ~/project_dir/public_html CityServerImpl

```

The files generated by `rmic` (in this case) are: `CityServerImpl_Stub.class` & `CityServerImpl_Skel.class`.

Polymorphic distributed computing is the ability to recognise (at runtime) the actual implementation type of a particular interface. We will use the example of a remote object that is used to compute arbitrary tasks:

- Client sends task object to compute server.
- Compute server runs task and returns result.
- RMI loads task code dynamically in the server.

This example shows polymorphism on the server, but it also works on the client, for example the server returns a particular interface implementation.

Our example task will be a simple interface that defines an arbitrary task to compute:

```

1 public interface Task extends Serializable
2 {
3     Object run();
4 }

```

Listing 4: Simple Task interface

We will also define a Remote interface:

```

1 import java.rmi.*;
2
3 public interface Compute extends Remote
4 {
5     Object runTask(Task t) throws RemoteException;
6 }

```

Listing 5: Simple Task interface

A task may create a Remote object on the server and return a reference to that object; the Remote object will be garbage-collected when the returned reference is dropped (assuming that no-one else is given a copy of the reference). A task may create a Serializable object and return a copy of that object; the original object will be locally garbage-collected when the Task ends. If the Task creates an object that is neither a Remote nor a Serializable object, a marshalling exception will be thrown.

```

1 import java.rmi.*;
2 import java.rmi.server.*;
3
4 public class ComputeServer extends UnicastRemoteObject implements Compute
5 {
6     public ComputeServer() throws RemoteException {}
7
8     public Object runTask(Task t)
9     {
10         return t.run();
11     }
12 }

```

Listing 6: Compute server implementation

```

1  public static void main(String args[])
2  {
3      System.setSecurityManager(new RMISecurityManager());
4      try
5      {
6          ComputeServer cs = new ComputeServer();
7          Naming.rebind("Computer", cs);
8      }
9      catch (Exception e)
10     {
11         // Exception handling
12     }
13 }

```

Listing 7: Compute server implementation

```

1  public class Pi implements Task
2  {
3      private int places;
4
5      public Pi (int places)
6      {
7          this.places = places;
8      }
9
10     public Object run()
11     {
12         // Compute Pi
13         return result;
14     }
15 }

```

Listing 8: Task to compute π

```

1  Compute comp = (Compute) Naming.Lookup("//www.t.nuigalway.ie/Computer");
2
3  Pi pi = new Pi(100);
4  Object piResult = comp.runTask(pi);
5
6  // print results

```

Listing 9: The client

In conclusion, RMI is flexible and allows us to pass objects (both Remote & Serializable) by exact type rather than declared type and download code to introduce extended functionality in both client & server. However, it is Java-only and has been superseded by SOAP & REST as the de-facto standards for communicating with remote services. Nonetheless, RMI is still worth learning to help understand concepts around distributed objects & distributed systems architecture.

3 Enterprise Java Beans

3.1 Distributed System Scenario

Imagine a worldwide financial company with 10,000 online customers that wants to add a new currency converter software component that is heavily used with 1,000 hits/second. The design will consist of the business logic and the distributed infrastructure. The distributed infrastructure includes security, load-balancing, transaction management, & object-relational mapping; Enterprise Java Beans takes care of this, and provides an API & framework.




Figure 3: Business logic, distribute the object, add security manager, add load balancing agent.

3.2 EJB

Enterprise Java Beans (EJB) is a server-side component architecture that enables and simplifies the process of building enterprise-class distributed object applications in Java. It allows you to write scalable, reliable, and secure applications without writing your own complex distributed object frameworks. EJB is a *specification*.




Figure 4: The EJB process

The **EJB Container** is where the EJBs run and is responsible for managing EJBs. The **EJB Server** is a runtime environment for container(s) that manages the low-level system resources.




Figure 5: The EJB server & containers




Figure 6: EJ Bean types

Session beans are “business process objects” (e.g., price quoting, order entry, video compression, stock trades, etc.) and live for as long as the client’s session. They are usable by 1 client at a time and are *not* shared. The EJB server manages the lifetime of beans. **Stateless session beans** are single request with no state kept, e.g., currency converter, compression utility, or credit card verification.

Entity beans / JPA represent persistent data. They are the object-oriented in-memory view of data in an underlying data store. They are long-lasting and have shared access. Sub-types of entity beans include: bean-managed persisted entity beans and container-manager persistent entity beans. **Bean-managed persistence** must be persisted manually and must look after saving, loading, & finding. They make use of a persistence API such as JDBC or SQL/J. **Container-managed persistence** is automatic persistence wherein the container/server looks after the loading, saving, & finding of component data. You must describe what you want persisted. Deployment tools provide support for defining simple object-relational mappings.

The client never invokes the bean instance, instead it invokes the **EJB object** by an invocation that is intercepted by the container, delegated to the bean instance. The EJB object is a surrogate, network-aware wrapper object that serves as a layer of indirection between the client & the bean; it is essentially the glue between the client & the bean. EJB objects must clone every business method that your bean class exposes, specified in the remote interface. All remote interfaces derive from `javax.ejb.EJBObject`.




Figure 7: EJB Objects

The **session bean interface** is implemented by all session beans and specifies lifecycle methods that may be implemented in the bean such as `setSessionContext`, `ejbCreate`, `ejbRemove`, `ejbPassivate`, & `ejbActivate`.

The **Java Naming & Directory Interface** is used to find an object. The resource (e.g., a bean) is associated with a nickname when deploying; clients of this bean can then use this nickname to look up the resource across a deployment. The client code looks up the reference in JNDI and calls business methods on the EJB object.

3.2.1 Entity EJBs

Entity EJBs are object-based representations of information-tier data such as data stored in a relational database. They represent a particular unit of data, e.g., a record in a database. There are two types of entity EJB:

- Bean-managed persistence;
- Container-managed persistence.

4 NodeJS

NodeJS is a JavaScript runtime environment that runs Google Chrome's V8 engine. It is a server-side solution for JavaScript which compiles JavaScript, making it quite fast. It was created in 2009 and designed for high concurrency, without threads or new processes. It has evented I/O for JavaScript, and never blocks, not even for I/O. Its goal is to provide an easy way to build scalable network programs. It provides a JavaScript API to access the network & file system and instead of threads, node uses an event loop with a stack which alleviates the overhead of context switching.

- JavaScript on the server-side ensures that communication between the client and the server will happen in the same language, with native JSON objects on both sides.
- Servers are normally thread-based, but Node is **event-based**; Node serves each request in an evented loop that can handle simultaneous requests.
- Node is a platform for writing JavaScript applications outside web browser, and is therefore not quite the same as the JavaScript we are familiar with in web browser: there is no DOM built-in to Node, nor any other browser capability.
- Node doesn't run in a GUI, but runs in the terminal or as a background process.

Threads	Event-Driven
Lock application / request with listener-workers threads.	Only one thread, which repeatedly fetches an event.
Uses incoming-request model.	Uses queue and then process it.
Multi-threaded server might block the request which might involve multiple events.	Manually saves the state and then goes on to process the next event.
Uses context switching.	No contention and no context switches.
Uses multi-threading environments where the listener & worker threads are used frequently to take an incoming-request lock	Uses asynchronous I/O facilities (callbacks, nor poll/select or O_NONBLOCK environments).

Table 1: Threads versus asynchronous event-driven

Ordinarily, a webserver waits for server-side I/O operations to complete while processing a web client request, thus **blocking** the next request to be processed. Servers generally do nothing but I/O, and scripts waiting on I/O requests degrades performance. Node processes each request as an event, it doesn't wait for the I/O operation to complete, making it **non-blocking**; it can therefore handle other requests at the same time. When the I/O operation of the first request is completed, it will callback the server to complete the request. To avoid blocking, Node makes use of the event-driven nature of JavaScript by attaching callbacks to I/O requests. Scripts waiting on I/O waste no space because they get popped off the stack when their non-I/O related code finishes executing.

4.1 MEAN

MEAN is a full stack solution consisting of MongoDB, Express, Angular, & node.




Figure 8: MEAN stack

5 Proxmox Virtualisation Environment

Proxmox is an open-source hyper-converged virtualisation environment. It has a bare-metal installer, a web-based remote management GUI, a HA cluster stack, unified cluster storage, and a flexible network setup. It has commercial support packages available at a reasonable cost. Proxmox uses the following underlying technologies:

- KVM (type 1 hypervisor module for Linux).
- QEMU hardware emulation.
- LXC Linux containers.
- Ceph replicated storage.
- Corosync cluster engine.

5.1 KVM

Kernel-based Virtual Machine (KVM) is a virtualisation infrastructure for the Linux kernel that turns it into a hypervisor. KVM requires a processor with hardware virtualisation extensions and a wide variety of guest operating systems work with KVM. It supports a paravirtual Ethernet card, a paravirtual disk I/O controller using VitrIO, a balloon device for adjusting guest memory usage, and a VGA graphics interface.

5.2 QEMU

QEMU (Quick Emulator) is an open-source hosted hypervisor that performs hardware virtualisation. It emulates CPUs through dynamic binary translation and provides a set of device models, enabling it to run a variety of unmodified guest operating systems. It uses KVM Hosting mode in Proxmox where QEMU deals with the setting-up and migration of KVM images. It is still involved in the emulation of hardware, but the execution of the guest is done by the KVM as requested by QEMU. It uses the KVM to run virtual machines at near-native speed (requiring hardware virtualisation extensions on x86 machines). When the target architecture is the same as the host architecture, QEMU can make use of KVM particular features, such as acceleration.

5.3 LXC

LXC (Linux Containers) is an operating-system-level virtualisation method for running multiple isolated Linux systems (containers) on a control host using a single Linux kernel. The Linux kernel provides the cgroups (control groups) functionality that allows limitation & prioritisation of resources (CPU, memory, block I/O, network, etc.) without the need for starting any virtual machines. It provides namespace isolation functionality that allows complete isolation of an application's view of the operating environment, including process trees, networking, user IDs, and mounted file systems. LXC combines the kernel's cgroups and support for isolated namespaces to provide an isolated environment for applications. Docker can also use LXC as one of its execution drivers, enabling image management and providing deployment services.

5.4 Ceph

Ceph is a storage platform that implements object storage on a single distributed computer cluster, and provides interfaces for object-level, block-level, & file-level storage. Ceph aims for completely distributed operation without a single point of failure, scalable to the exabyte level. Ceph's software libraries provide client applications with direct access to the Reliable Autonomic Distributed Object Store (RADOS) object-based storage system. Ceph replicates data and makes it fault-tolerant, using commodity hardware and requiring no specific hardware support. As a result of its design, the system is both self-healing and self-managing, aiming to minimise administration time and other costs. When an application writes data to Ceph using a block device, Ceph automatically striped and replicates the data across the cluster. It works well with the KVM.




Figure 9: Ceph architecture




Figure 10: Ceph internal organisation

5.4.1 Ceph Network

To create a Ceph ring₀ network, each node must be reachable on `rin0`. The firewalls on each node will need to be checked to verify this. Proxmox distribute their own Ceph package as of version 5.1:

- `pveceph install` will install the latest stable repositories & packages – must be run on each node individually.
- `ceph init --network x.x.x.x/y` must be run on the first node only.
- `ceph createmon` must be ran on each node.

5.4.2 Ceph OSDs

We can add disks as **Object Storage Devices (OSD)s** on each node. The accurate network time is also very important to avoid “clock skew”. The latest network time system daemon `system.time?` is much better than `ntpd`. QEMU has a new time source driver which can be run in guests needing accurate time.

Pools are individual storage blocks:

- `size` is the number of replications (OSDs) per block.
- `min-size` is the minimum number of OSDs (replications) each block must be on to allow read-write status.
- `add-storage` option automatically adds the storage block to the hosts rather than having to manually copy the Ceph keys to each host to allocate the storage.

The client (Proxmox) interacts with one OSD only. This OSD then write to and confirms write on each OSD in the block before confirming write completion. Writes are actually made to the journal rather than the block level device for speed. This primary OSD manages all interactions with both the client and the replication OSDs. In case the primary manager is lost, a backup OSD will take over as primary.

Crush maps define the actual storage blocks (these are very complicated so don't change the default settings!). As new OSDs are added, Ceph will attempt to re-allocate data across blocks to improve access & availability. If an OSD gets removed, Ceph will rebalance data once the OSD is marked as OUT (300 seconds by default). Use `ceph noout` to avoid rebalancing, e.g., for maintenance.

5.5 VM Installation

5.5.1 Hard Drives

For hard drives, `virtio-scsi` and `scsi` are the best-performance options. On Windows VMs, this can be a chore as it is necessary to use a second boot CD to install `virtio` drivers. `No-cache` is the best compromise option for local disks, as the write-back is the fastest, although it is unsafe on Ceph.

5.5.2 Memory

Fixed allocation with ballooning is the best way to allocate RAM. Over-provisioning is possible but dangerous as guests may crash if RAM is not available. Auto-allocation of memory means that required RAM may take up to 30 seconds to be available; it is best to leave swap enabled as this way, swap is a last option before crashing.

5.5.3 VM Backups

There are two types of backup types:

- **Snapshot** leaves the guest running and intercepts all write operations, writes them to the backup if the block is already backed up, then to the guest. This slows the guest I/O down to the speed of the backup medium.
- **Stop** causes the guest to shutdown, then restarts and does backup before making the guest available.

5.5.4 VM Migration

For guests on local storage, migration must done offline. Any storage used in the guest (e.g., ZFS) must be available on the target node. guests can be live moved to shared storage (e.g., NFS or Ceph) and then live migrated.

5.5.5 VM Cloning

Linked clones allow the fast spin-up of machines as only diverging blocks need to be written to the disk. Linked clones require file-level storage system, i.e., snapshot-able storage. The conversion of a VM to a template sets the image as read-only.

5.5.6 VM Imports

To perform **OVA import**, first unpack the OVA, for example onto a NAS. Then run `qm help importovf` for details of the import command. To perform **disk import**, run `qm help importdisk`. `vmdebootstrap` can be used to build Debian disk images programmatically. `qm help create` can be ran for details on creating VMs programmatically.

Note that Windows disk images will not have any `virtio` drivers installed by default: the hard disk types must be SATA, the network devices must be E1000. `Spice-space spice-guest-tools` can be used to install all `virtio` drivers into Windows images. The Spice repository on GitHub has the source code for the installation tools.

5.5.7 User Authentication

PAM authentication can be used for per-machine authentication (may be possible to integrate radius). **Proxmox authentication server** replicates authentication across all nodes.

5.6 Proxmox Cluster

The **Proxmox VE cluster manager pvecm** is a tool to create a group of physical servers called a **cluster**. It uses the Corosync Cluster Engine for reliable group communication, and such clusters can consist of up to 32 physical nodes or more, dependent on the network latency (must be less than 2 milliseconds). `pvecm` can be used to create a new cluster, join nodes to a cluster, leave the cluster, get status information, and do various other cluster-related tasks.

Grouping Proxmox hosts into a cluster has the following advantages:

- Centralised, web-based management of a multi-master cluster: each node can do all management tasks.
- `pmxcfs`: a database-driven file system for storing configuration files, replicated in real-time on all nodes using the corosync cluster engine.
- Migration of VMs & containers between physical hosts.
- Fast deployment & cluster-wide services like firewall and High Availability (HA).

5.7 High Availability

Items managed under HA are referred to as *resources*:

- The HA cluster is managed by `pve-ha-crm.service`.
- The local HA resources are managed by `pve-ha-lrm.service`.

The Guest HA is managed either through the dropdown on the guest window, or HA options on the Datacenter and this allows a guest VM to be automatically migrated or restarted on a different node if it is detected as down, e.g., because of node failure or maintenance.

Ensure that `pve-ha-crm` and `pve-ha-lrm` are both running under node `-> services`. All migrations and other actions on HA resources are managed by the HA daemon. The task viewer only shows status of the request to HA daemon to carry out the task, not of the actual task.

Migrations (generally, but particularly under HA conditions) may fail due to a number of causes, including:

- The guest has local attached storage which is not available on the target node.
- The guest has NUMA (Non-Uniform Memory Access) or other CPU settings not present on the target node.

Changing the HA manager state for a VM will cause the VM state to change. If any node hosting a HA resource loses corosync quorum:

1. The `pve-ha-lrm.service` will no longer be able to write to the watchdog timer service.
2. After 60 seconds, the node will reboot.
3. After a further 60 seconds, the VM will be brought up on a different node.

5.7.1 HA Groups

Group members will prefer selected nodes if available:

- If `restricted` is selected, members will only run on selected nodes.
- Guests will be stopped by the HA manager if the node(s) become unavailable.
- If `nofailback` is not selected, guests will try to migrate back to a preferred node once it becomes available again.

5.8 Performance Benchmarking

- `iperf` to test network throughput.
- `sysstat` to monitor system statistics.
- `iostat` to test I/O throughput.

6 Cloud Computing

Cloud computing solves web-scale problems, uses large data centers, typically uses different models of computing, and produces highly-interactive web applications. **Web-scale problems** are typically highly data-intensive and may also be processing intensive.