

Introduction

COMP750 Distributed Systems

What is a Distributed System?

A distributed system consists of a collection of autonomous computers linked by a computer network and equipped with distributed system software.

What is a Distributed System?

A distributed system is a collection of independent computers that appears to its users as a single coherent system.

What is a Distributed System?

“A distributed system is one where you can't get your work done because a computer you didn't even know existed has crashed.”

- Leslie Lamport

Almost Distributed Systems

- Multiprocessor
- Server with multiple terminals

Why distributed systems?

- People are distributed, data are distributed
- performance / cost
- Modularity
- Expandability
- Availability
- Scalability
- Reliability

Goals


- Connect users to resources
- Hide the fact that resources are distributed
- Open
- Scalable
- Robust
- Efficient and powerful

Transparency in a Distributed System

Transparency	Description
Access	Hide differences in data representation and how a resource is accessed
Location	Hide where a resource is located
Migration	Hide that a resource may move to another location
Relocation	Hide that a resource may be moved to another location while in use
Replication	Hide that a resource may be shared by several competitive users
Concurrency	Hide that a resource may be shared by several competitive users
Failure	Hide the failure and recovery of a resource
Persistence	Hide whether a (software) resource is in memory or on disk

Different forms of transparency in a distributed system.

Structure of a Distributed System



A distributed system organized as middleware.
Note that the middleware layer extends over multiple machines.