Lecture 6
Distributed Computing

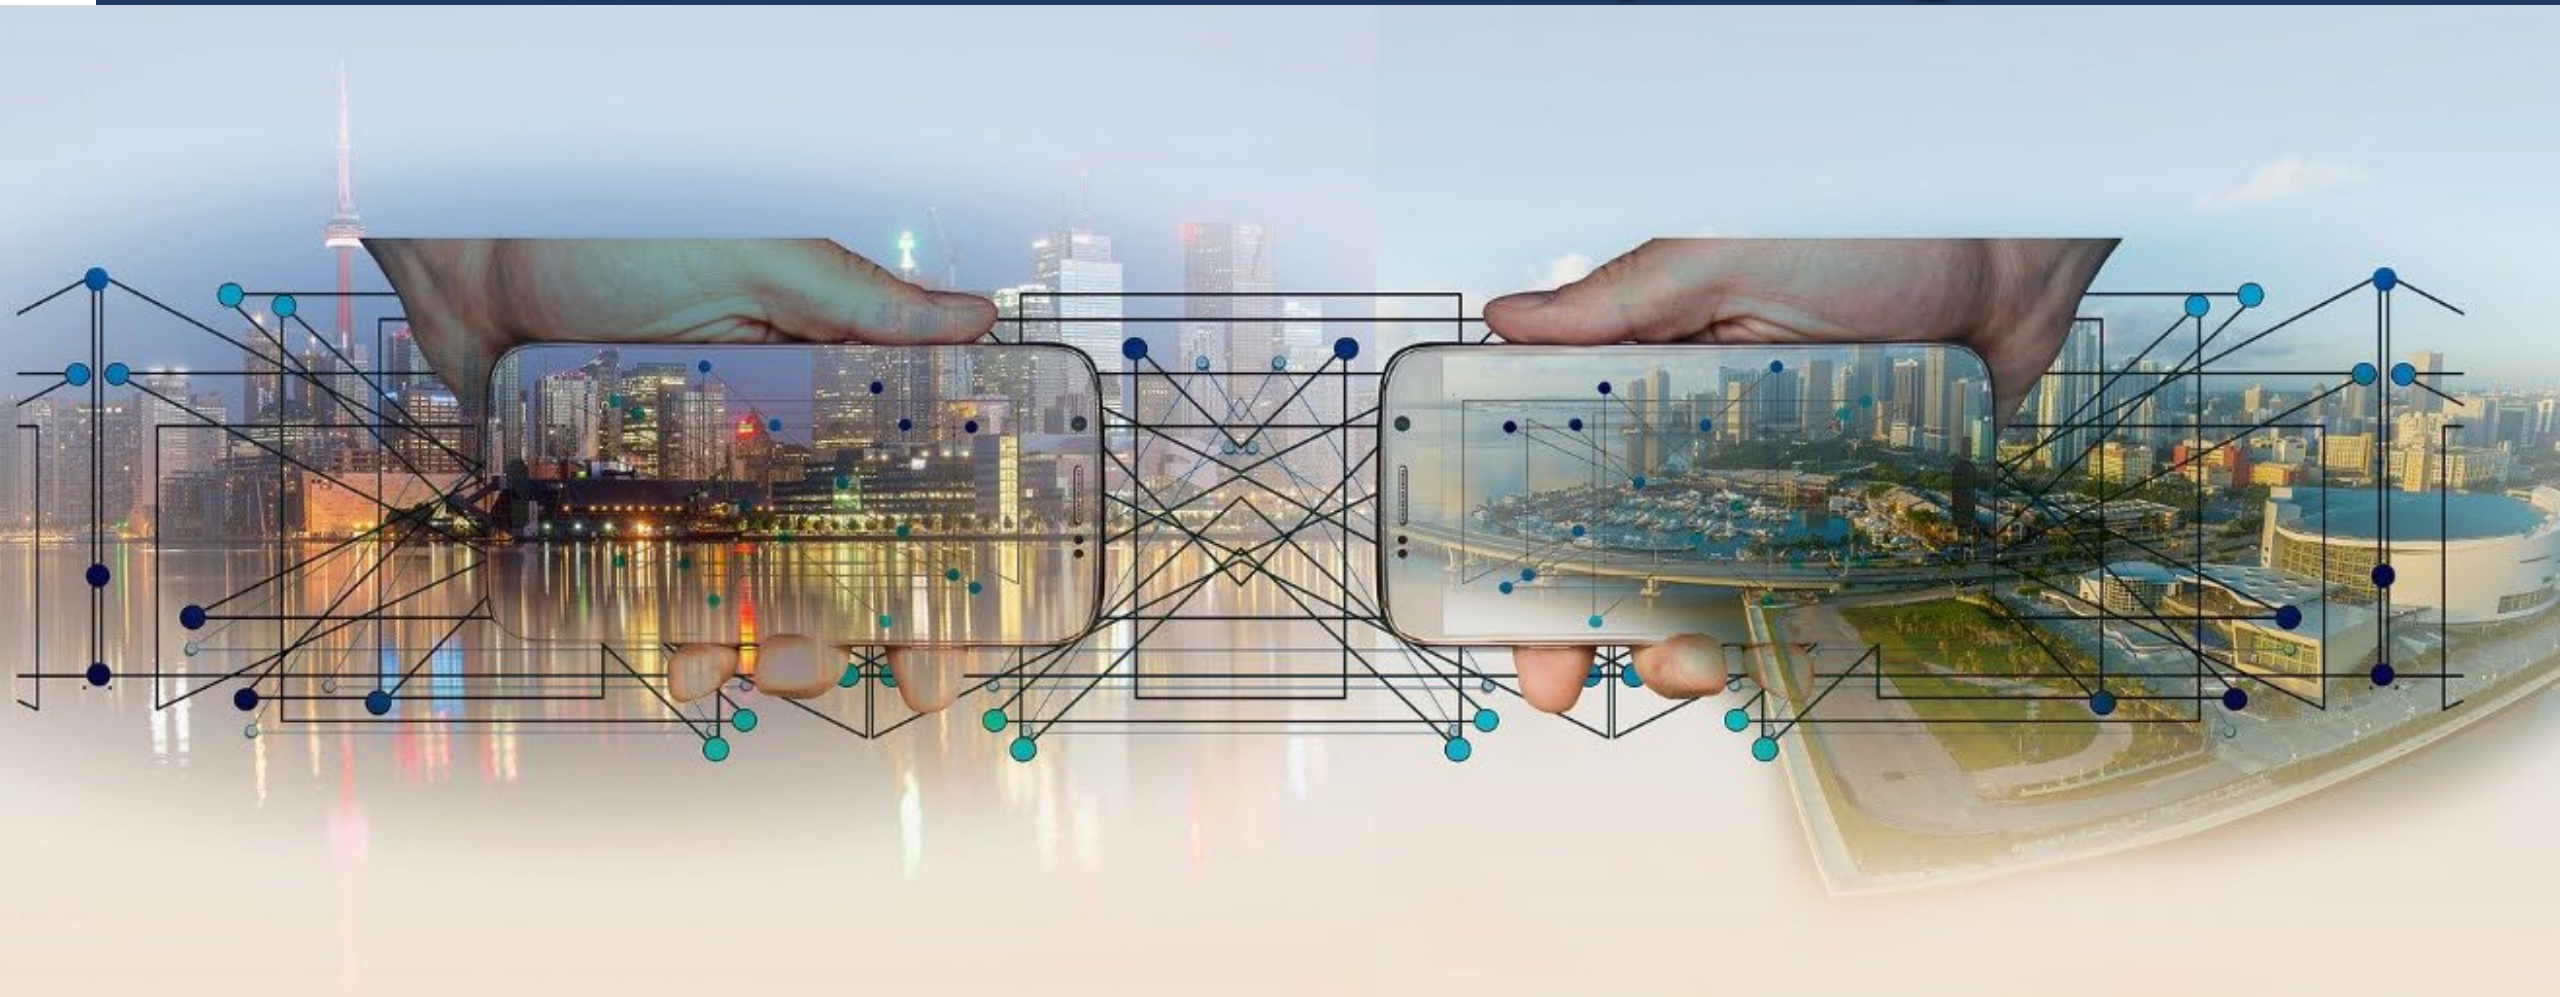

Why distributed systems ?

- **Fact**: Processor population is exploding. Technology has dramatically reduced the price of processors.
- Geographic distribution of processes
- Resource sharing as used in P2P networks
- Computation speed up (as in a grid)
- Fault tolerance

What is Distributed Systems?

Hype Cycle
(https://www.gartner.com/)

Each Hype Cycle drills down into the five key phases of a technology's life cycle.

- **Innovation Trigger:**
  - A potential technology breakthrough kicks things off. Early proof-of-concept stories and media interest trigger significant publicity. Often no usable products exist and commercial viability is unproven.
- **Peak of Inflated Expectations:**
  - Early publicity produces a number of success stories — often accompanied by scores of failures. Some companies take action; many do not.
- **Trough of Disillusionment:**
  - Interest wanes as experiments and implementations fail to deliver. Producers of the technology shake out or fail. Investments continue only if the surviving providers improve their products to the satisfaction of early adopters.
- **Slope of Enlightenment:**
  - More instances of how the technology can benefit the enterprise start to crystallize and become more widely understood. Second- and third-generation products appear from technology providers. More enterprises fund pilots; conservative companies remain cautious.
- **Plateau of Productivity:**
  - Mainstream adoption starts to take off. Criteria for assessing provider viability are more clearly defined. The technology's broad market applicability and relevance are clearly paying off.

Hype Cycle for Emerging Technologies 2024

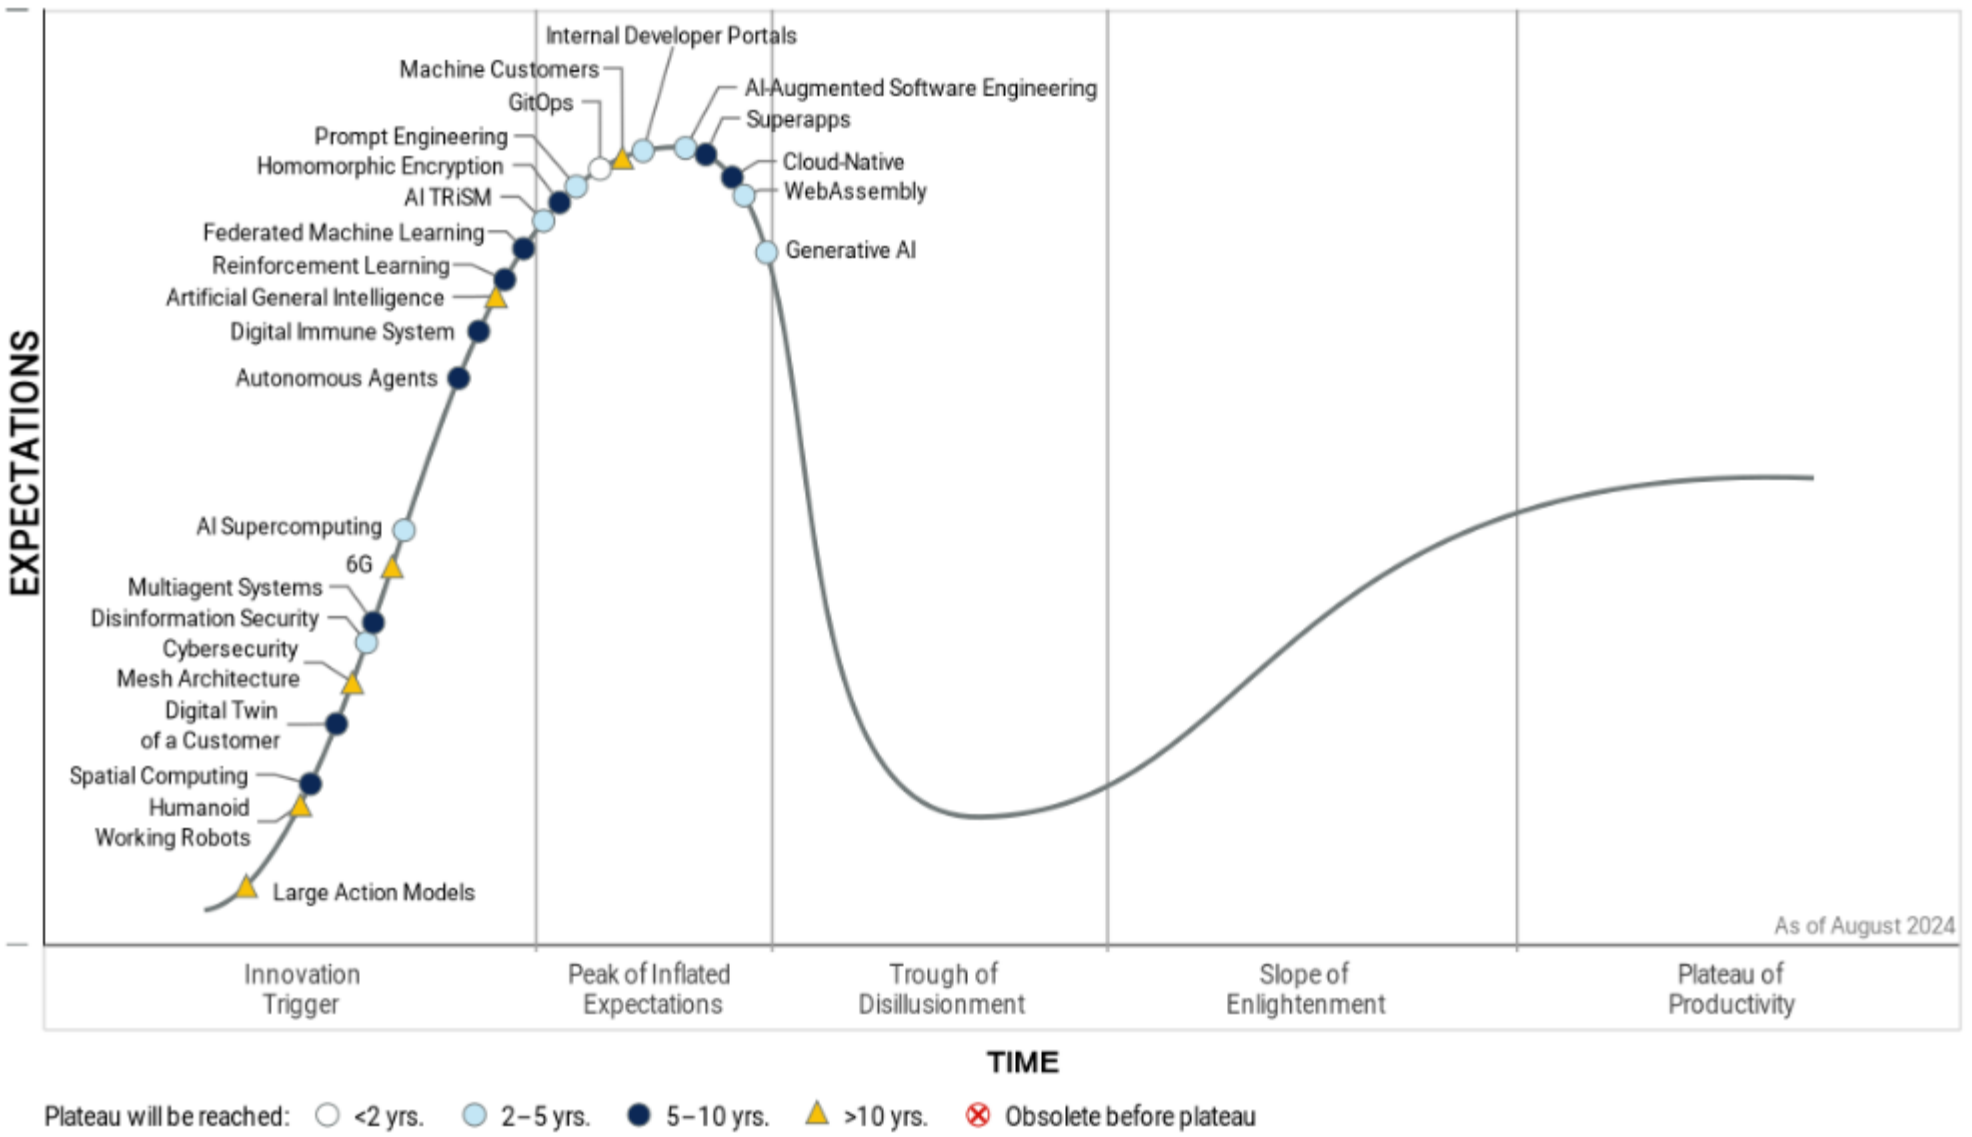

Gartner

Hype Cycle for Emerging Technologies 2023

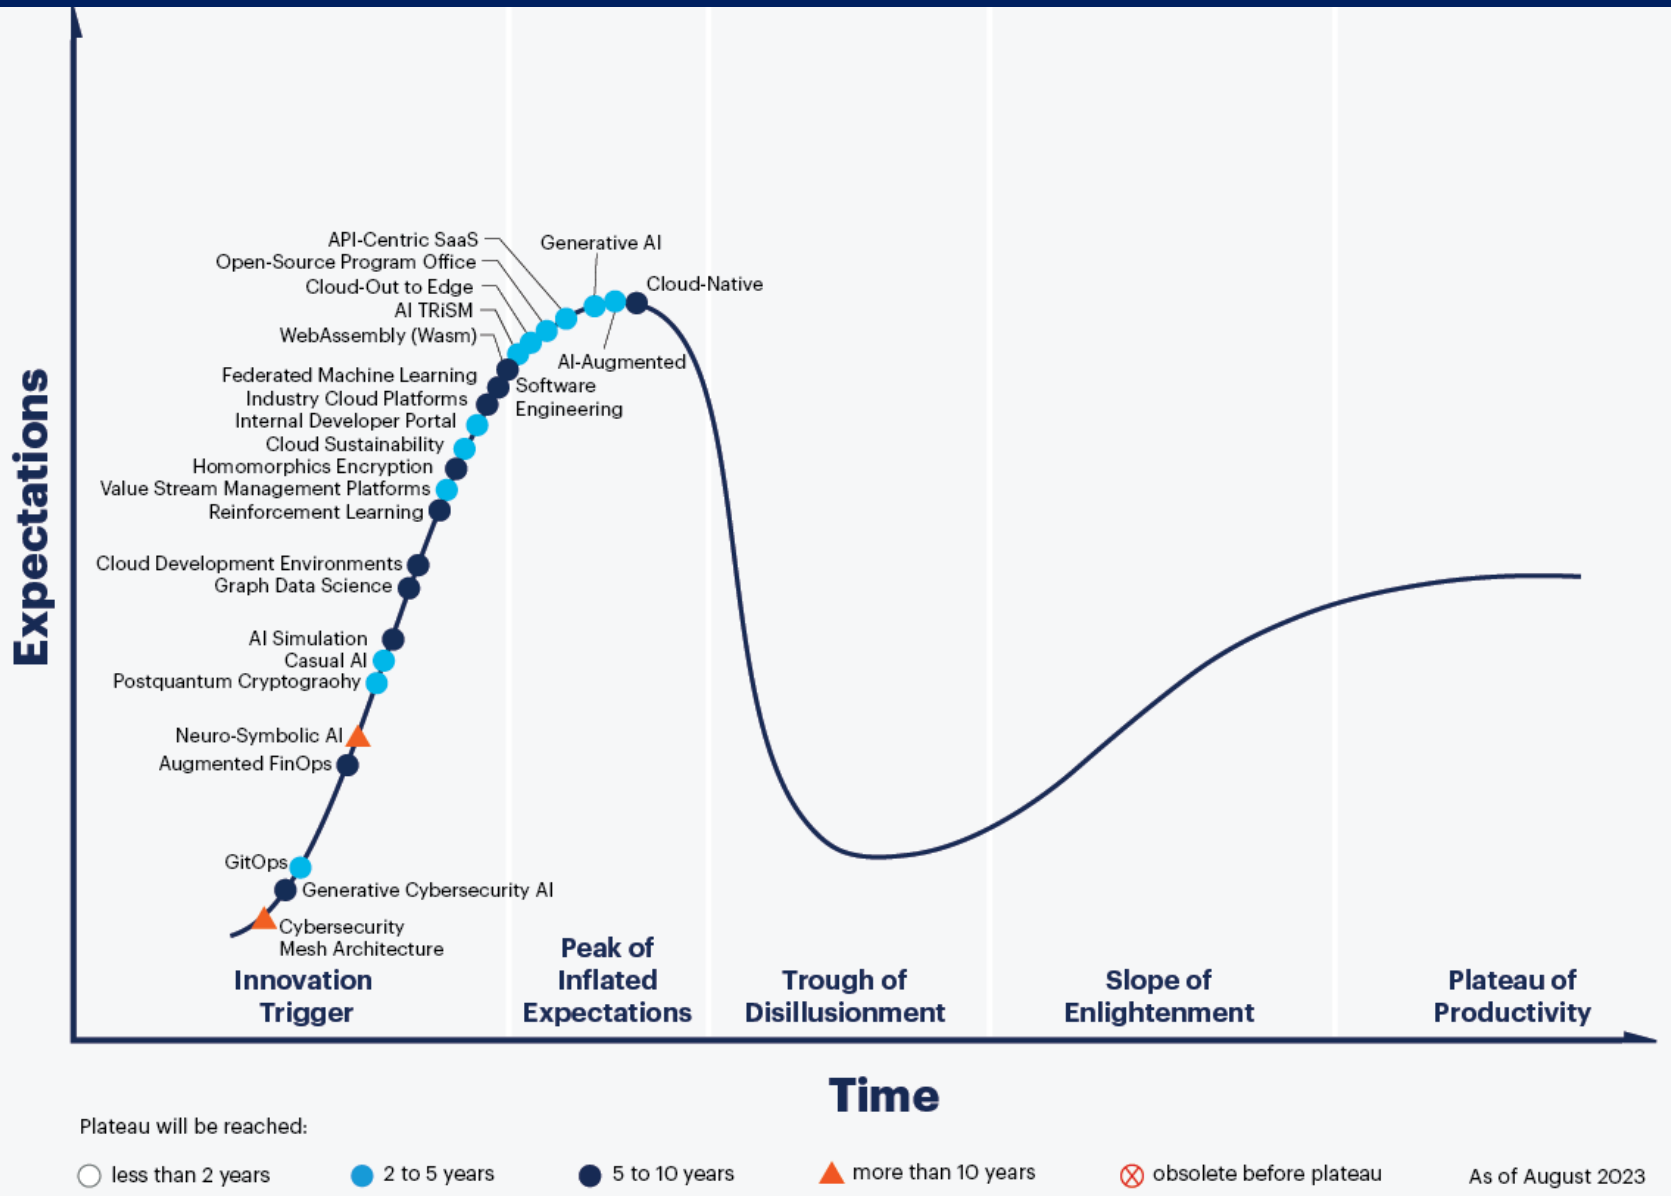

gartner.com

What is a distributed system? (1)

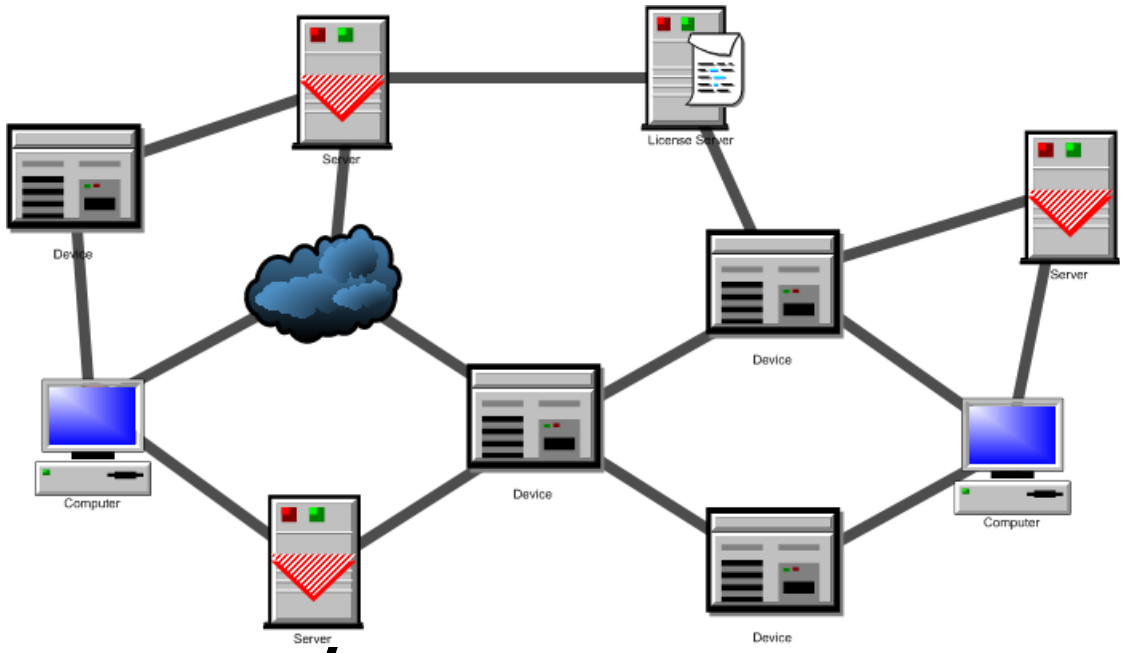

- A network of processes/resources.
- The nodes are processes/resources, and the edges are communication channels.

What is a distributed system? (2)

- The logical distribution of functional capabilities
  - Multiple processes
  - Interprocess communication
  - Disjoint address space
  - Collective goal

Collective Goal?

- Don't hold your breath:
  - Biocomputing
    - Using biologically derived molecules
  - Nanocomputing:
    - the manipulation of matter on an atomic and molecular scale
  - Quantum computing
    - use of quantum-mechanical phenomena, such as superposition and entanglement
  - …
- It all boils down to…
  - Divide-and-conquer
  - Throwing more hardware at the problem

Divide and Conquer

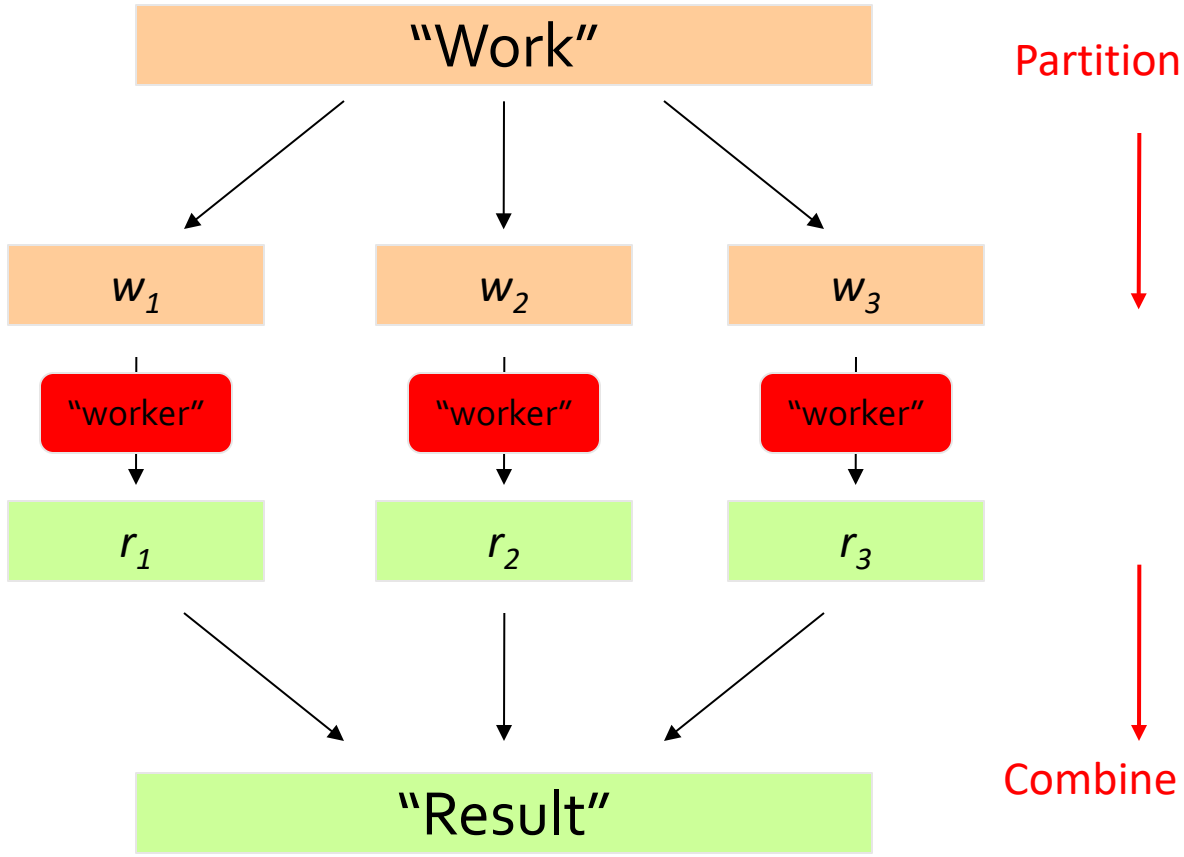

Reference: Lecture Note by Jimmy Lin, Univ. of Maryland

What is a distributed system? (2)

Different Workers

- Different threads in the same core
- Different cores in the same CPU
- Different CPUs in a multi-processor system
- Different machines in a distributed system

Reference: Lecture Note by Jimmy Lin, Univ. of Maryland

What is a distributed system? (2)

Choices, Choices, Choices

- Commodity vs. “exotic” hardware
- Number of machines vs. processor vs. cores
- Memory vs. disk vs. network bandwidth
- Different programming models

Reference: Lecture Note by Jimmy Lin, Univ. of Maryland

A classification

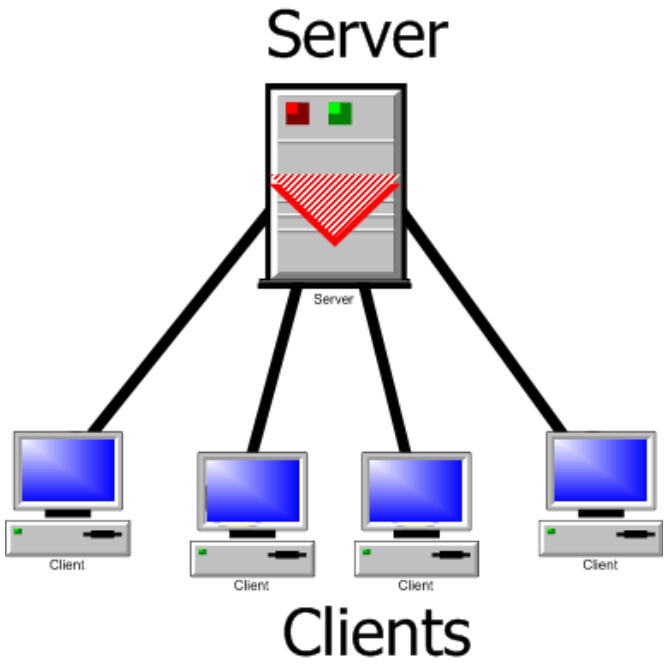

Client-server model

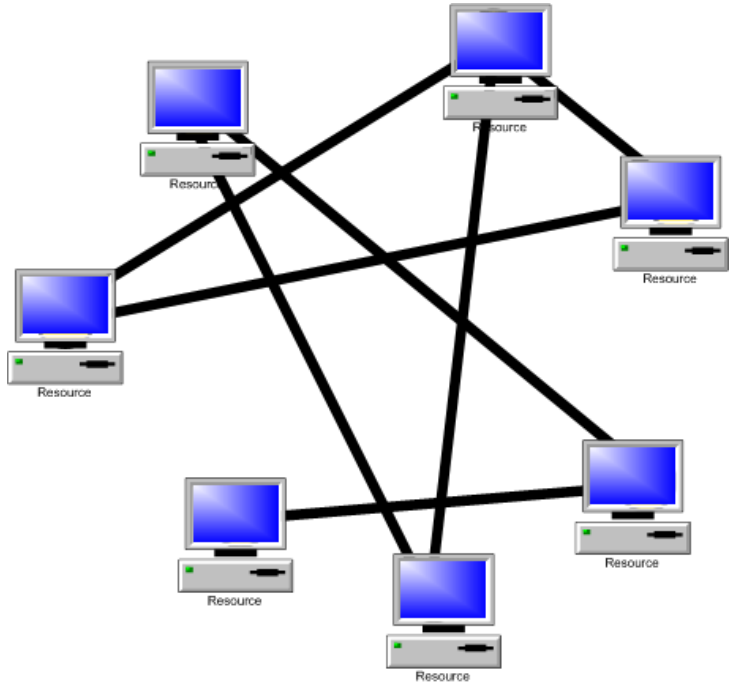

Peer-to-peer model

Parallel vs. Distributed

- In both parallel and distributed systems, the events are *partially ordered*.

- In parallel systems, the primarily issue is speed-up
- In distributed systems the primary issues are fault-tolerance and availability of services

Possible Services

- Internet banking
- Web search
- Net meeting
- Distance education
- Internet auction
- Google earth
- Google sky
- And so on…

More Examples

- Large networks are very commonplace these days. Think of the world wide web. Other examples are:
  - Ubiquitous Computing
  - Cloud computing
  - Grid computing, Grid computing networks
    - Ex. Computational grids (OSG (Open Science Grid), Teragrid, SETI@home)
  - Sensor networks
  - Network of mobile robots
  - And so on…

Sensor Network

Checking the structural integrity

Mobile robots

Cloud Computing

Image Source: www.vemurivenkatrao.com/nature/cloud/

Important issues

- Knowledge is local
- Clocks are not synchronized
- No shared address space
- Topology and routing
- Scalability
- Fault tolerance

Some common sub-problems

- Leader election
- Mutual exclusion
- Time synchronization
- Distributed snapshot
- Replica management