

Mainframe versus Centralized versus Distributed System

How the Terms “Mainframe” and “Distributed” Are Used in This Course

The terms “mainframe” and “distributed” are used throughout this course in text, graphics, games, and video and audio recordings. Nearly all of the comparison details currently available define the distributed system as a platform that has the following characteristics:

- Assigns a single server to support a single application workload
- Is externally networked
- Has dedicated memory and I/O
- Supports numerically intensive applications
- Supports availability through failover

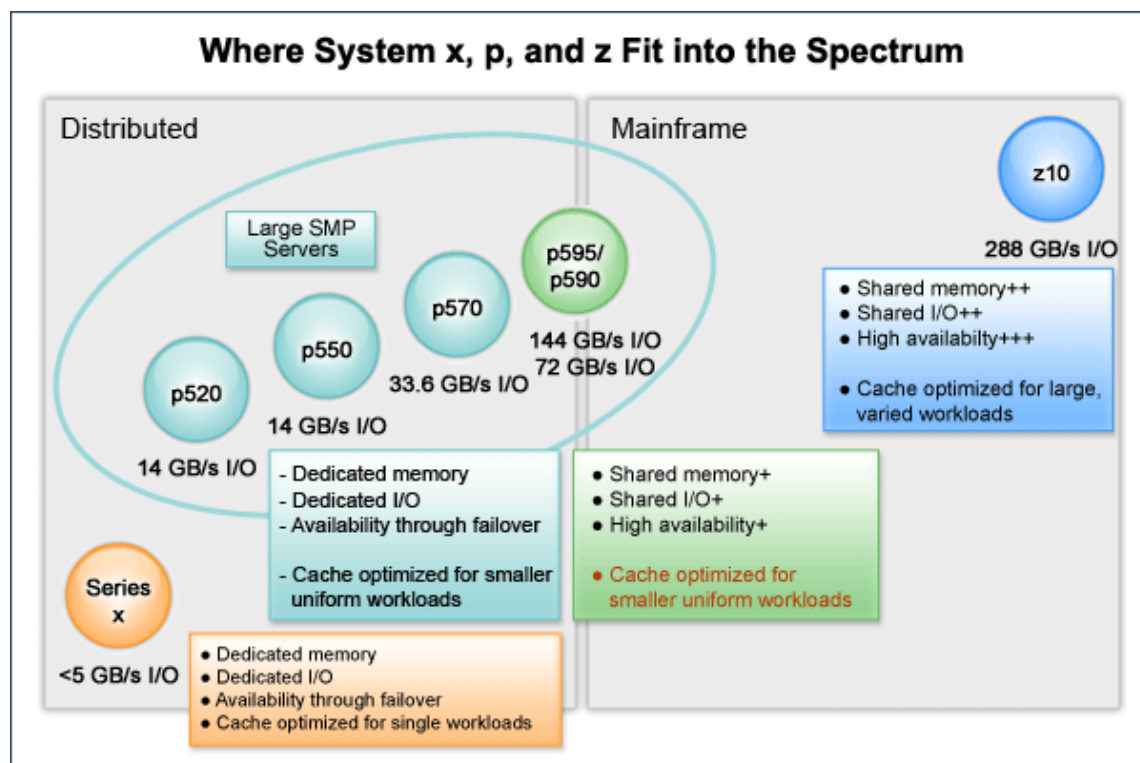
At IBM, the mainframe is generally referred to as System z, which includes z/OS, z/VM, and Linux on IBM System z. System z is designed with high Qualities of Service (QoS) standards. As such, the mainframe has the following characteristics:

- Provides high availability
- Runs multiple application workloads on a single image
- Supports highly data-intensive applications
- Has shared memory and I/O

These descriptions form the basis of how the terms “mainframe” and “distributed” are used throughout this course. Because distinctions and differences exist between the way z/OS, z/VM, and Linux on System z perform certain tasks, they are qualified separately in this course.

Mainframe versus Centralized versus Distributed

Due to upgrades in technology over the past few years, several IBM products are beginning to overlap the familiar, traditional boundaries of the mainframe and the distributed system, as depicted in the following graphic.



Mainframe versus Centralized versus Distributed System

Because of this overlap, the term “centralized” is beginning to replace the term “mainframe,” where centralized has the following characteristics:

- Runs general commercial workloads (low numerical intensity, high data intensity)
- Runs unpredictable heterogeneous workloads (mix of long and short transactions on the same data)
- Has no difficulty partitioning workloads (workloads that frequently shift data access patterns and processing patterns)
- Batches workloads on a common database (very I/O intense)
- Runs mission-critical workloads
- Generally has more than 16 cores, for example, Power 570, Power 595, System z9®, System z990, and System z10™
- Runs each virtual machine to scale dynamically to utilize the full capacity of the CPU (PowerVM™, zVM®)
- Provides utilization rates in the 70%+ range (System z rates are consistently in the 90%+ range)
- Offers highest business value and lowest total cost of ownership (TCO)

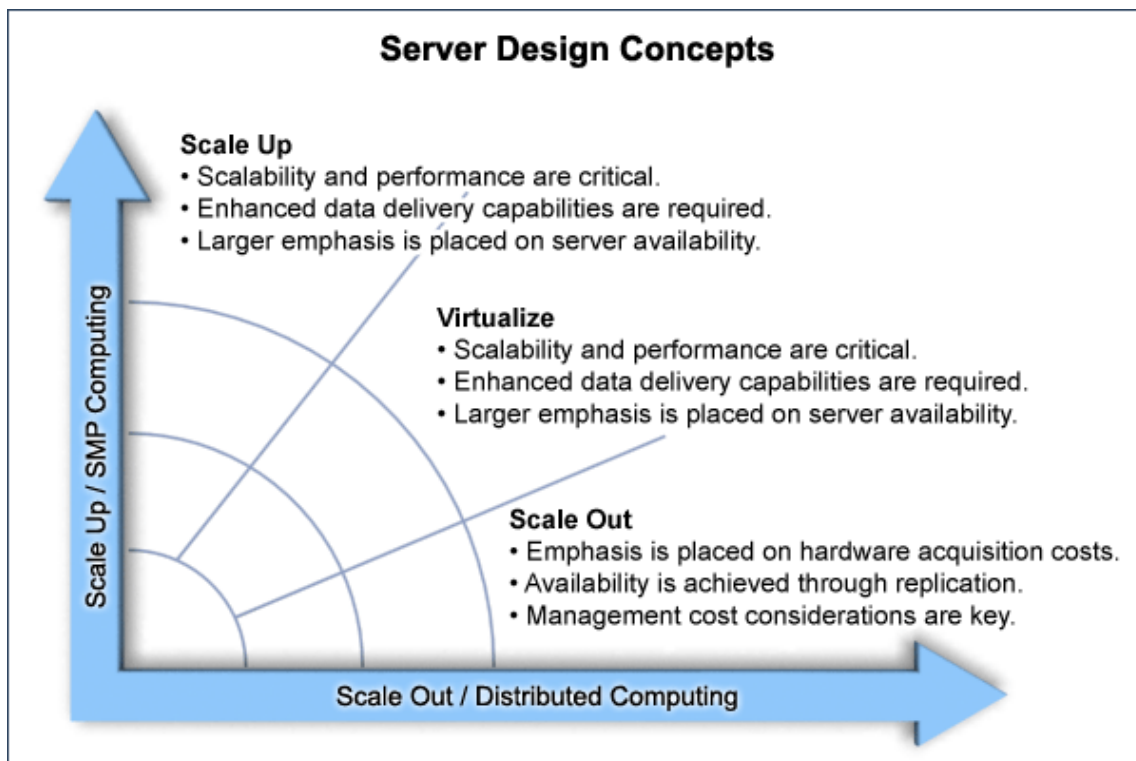
The term “distributed” remains in use, generally with the following characteristics:

- Runs specialized workloads (numerically intensive, computationally intensive), although the new chip speed on the System z10 is bringing the System z10 up to 85% of the high-end System p chip
- Runs “well-behaved” workloads (short, stateless transactions)
- Runs data-intensive workloads with low synchronization requirements
- Runs less mission-critical workloads
- Has virtual machines that are limited to eight or fewer cores:
 - Currently, all limited to four or fewer cores (VMWare, IVM, LDomS)
 - Two-chip scalability for dual core, one-chip scalability for quad core, for example, BladeCenter®, x86, Power Systems with fewer than 16 cores (Power 520, Power 550), SPARC (Sun/Fujitsu), Itanium (HP/SGI)
- Has hardware partitioning (DR, vPar, nPar)
- Typically provides average utilization rates in the 5% to 25% range
- Offers generally lower total cost of acquisition (TCA), but higher systems management and environmental impact, which contributes to a higher TCO

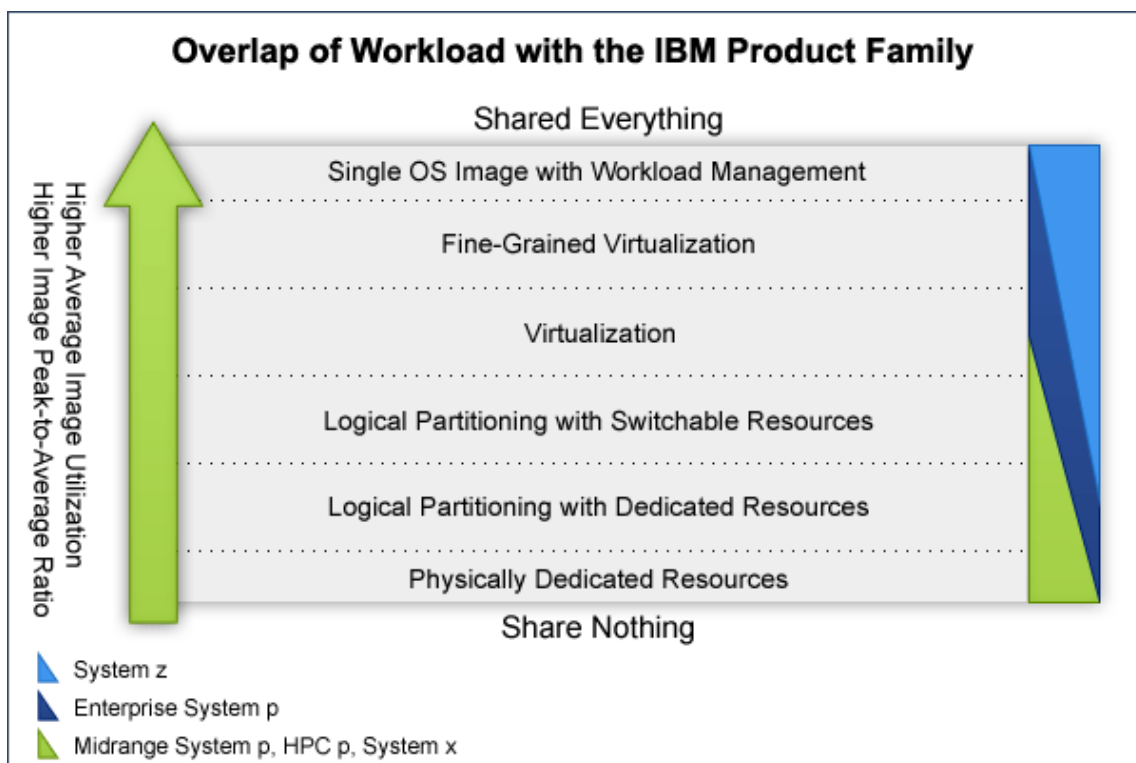
Centralized servers have shared resources, complete all their work in one place, and support high availability through failure avoidance and recovery. Distributed servers, which are often referred to as clustered servers, have dedicated resources, have applications in multiple servers externally networked together, and support high availability through redundancy and failover.

With these definitions in mind, System z can support both the scale-out characteristic of distributed systems with Linux on System z with z/VM and the scale-up characteristic of centralized servers with z/OS, as the following graphic illustrates.

Mainframe versus Centralized versus Distributed System



The following graphic pictorially shows the overlap of this workload with the IBM product family. Virtualization on the platforms also reveals more overlap between the product sets. An important point to note is that System z has included virtualization for many years, becoming highly refined and proven, while the distributed platform has only more recently adopted virtualization.



Mainframe versus Centralized versus Distributed System

Summary

The key information to remember is that customers are looking to IBM and new technology to help them move toward a more cost-effective IT environment that accommodates business growth and flexibility. Some traditional distributed platforms are yielding to the demands of a centralized infrastructure—one that historically has been reserved for the mainframe and System z. However, System z has been proven to be a cost-effective IT environment that now competes in processing speed, handling mixed workloads with a CPU utilization of 90%+. The strengths of System z directly meet customers' business requirements, and IBM currently sees that clients more frequently are selecting System z to meet those demands.

Regardless, it is critical that customers buy the right product to solve their business problems, and, in all cases, the goal is for IBM to provide that product. If you have any questions about which product best meets your customers' needs, be sure to involve the appropriate hardware and software experts from Systems and Technology Group (STG) and Software Group (SWG).