

Enterprise Computing Insights

What is Centralized Computing?



#2 IN A SERIES

Centralized computing refers to a model of computing used to conceptualize the relationship between end users and computers. When computers were first introduced into businesses, centralized computing was the only kind of computing since there was typically only one computer in a business. All users of the computer came to the (central) computer to get work done. Eventually, text based dumb terminals (with no local processing capability) were directly connected to the computer to allow users to interact with the central computer without physically going to the machine for service. Today, dumb terminals have been replaced with “thin clients” (with little or no local processing power) or “personal computers” (which do have local processing capability and a graphical user interface) to connect to a central computer over a network and thus interact with it. A working definition of centralized computing is: a small number of powerful servers at a common location that satisfy the computing needs of a large number of end users.

Over time computers became much more prevalent in business due to the reduction in their size, reduced need for a conditioned environment and greatly reduced acquisition cost. This led to computers being purchased and used at the “departmental” level within businesses. Departmental use of computers led to their proliferation across a business and therefore their geographic distribution across a business as well. Other business factors such as security, mergers, etc. also contributed to the geographic dispersal of computers within a business. Initially, these geographically dispersed computers were used as “stand alone” independent computer systems. The model of computing characterized by multiple computers that did not communicate or interact with each other is called “decentralized computing”. Each department that purchased its own computer was responsible for its management and control.

Over time these independent computers were connected using networking technology so that they could communicate and interact with each other. That is the type of geographically dispersed configurations we find in use today and it is known as “distributed computing”. There is usually no assumed homogeneity among the computers and the structure of the distributed configuration can and does change over time. The term “distributed computing” has continued to evolve and today has taken on an even broader meaning than just geographic distribution of interacting servers. For example, independent programs running on the same server and communicating with each other are sometimes characterized as “distributed computing” since there is a similarity to independent computers communicating with each other. A working definition of “distributed computing” is: several independent entities (usually computers) in multiple locations whose programs can communicate with each other, either synchronously or asynchronously, over a network using messages or other forms of interaction. The distributed configuration collectively satisfies the needs of a large number of end users and some or all of the individual computers in the configuration are used to satisfy the needs of local users.

If we examine each of these computing models further, we can identify advantages and disadvantages associated with each (with the advantages and disadvantages based on a comparison to each other). Think of an analogy between the central heating/cooling of a building and the distributed heating/cooling of a building and you can probably identify a number of advantages and disadvantages yourself.

Some of the advantages of a purely centralized computing configuration include:

- ease of management; there are relatively few (perhaps one) computers to manage
- enhanced security; the physical and logical securing of the computing environment can be more

easily managed since there is only one location and a few computers

- ease of control; the introduction of change can be managed closely since there is only one location and a few computers
- reduced cost of ownership; due to fewer computing elements to manage and therefore fewer people needed to manage them.
- multiple types of workload; all of the work associated with the business runs at the central computing location. A variety of work running on the server(s) means that the equipment can be kept very busy (high utilization) most of the time which is helpful in making the best economic use of the investment in the computing equipment.

Some of the disadvantages of a purely centralized computing configuration may include:

- reduced availability; if there is exactly one computer and it fails, all end users of the central computer no longer have their computing needs met. This was more of a problem in the early days of computing when there might have been only one computer to service the needs of a business. Today, centralized computing facilities take advantage of clustering and other techniques for high availability. In fact, these centralized facilities are often more highly available than distributed systems since the mean time before failure (MTBF) of the hardware is measured in decades and the availability of the centralized hardware/software configuration can approach 99.999%.
- Perceived bureaucracy; change is usually introduced in a centralized configuration in a very deliberate and controlled fashion. This is necessary to avoid introducing a change that causes a failure or an outage. This tight control over changes can be perceived as making the computing facility inflexible and rigid when it comes to introducing changes.

Some of the advantages of a purely distributed computing configuration include:

- improved availability; since there are multiple computers in the distributed configuration, a single failure usually affects only some of the end users, not all of them. This assumes that the failed

computer does not manage or own a resource that is critical to all users of the distributed configuration.

- local control of change management; this makes it easier to make changes since the local management understands their environment and configuration very well. There will likely be fewer reviews and approvals needed to make the change. And if something should go wrong when the change is introduced, only the local computing will likely be affected; the rest of the distributed computers and end users will not likely be affected.

- ease of adding additional capacity; growth can occur by simply adding additional distributed computers into the configuration instead of having to upgrade a central computer. Incremental growth may be available in smaller increments than that available on a large central server.

Some of the disadvantages of a purely distributed computing configuration may include:

- Increased number of personnel needed; more time and effort (which implies cost) will likely need to be dedicated to managing all of the distributed computers. Skilled personnel need to be available locally to manage and operate the computing systems although some of this can be done remotely.
- More difficult to meet global business objectives/mandates (for security, auditing, etc.) since they need to be communicated and implemented by so many more people/computers. There are also more opportunities for errors in the implementations since more computers are involved.
- Growth can occur by simply adding additional distributed computers. Since acquisition costs can be modest, this may lead to uncontrolled growth.
- Local use and management of the compute resources may lead to them being used for only one purpose and therefore used inefficiently (e.g. low utilization).

While we can discuss each of these models in a purely idealized sense, in most real world situations, you will usually find them used in conjunction with each other – especially in a large organization. It is

not unusual to find a large number of distributed servers surrounding a small collection of Enterprise Servers in a modern data center¹. And many distributed servers scattered geographically throughout the organization. More often than not however, all of the equipment is the responsibility of the business' IT organization.

We will explore other aspects of Enterprise Computing in subsequent articles.

¹ See ECI No. 1