

# Physical Infrastructure and Costs of WSC

- To build a WSC, you first need to build a warehouse. location for a WSC means proximity to Internet backbone optical fibers, low cost of electricity, and low risk from environmental disasters, such as earthquakes, floods, and hurricanes. For a company with many WSCs, another concern is finding a place geographically near a current or future population of Internet users, so as to reduce latency over the Internet. There are also many more mundane concerns, such as property tax rates
- Infrastructure costs for power distribution and cooling dwarf the construction costs of a WSC, so we concentrate on the former. Figures below show the power distribution and cooling infrastructure within a WSC.

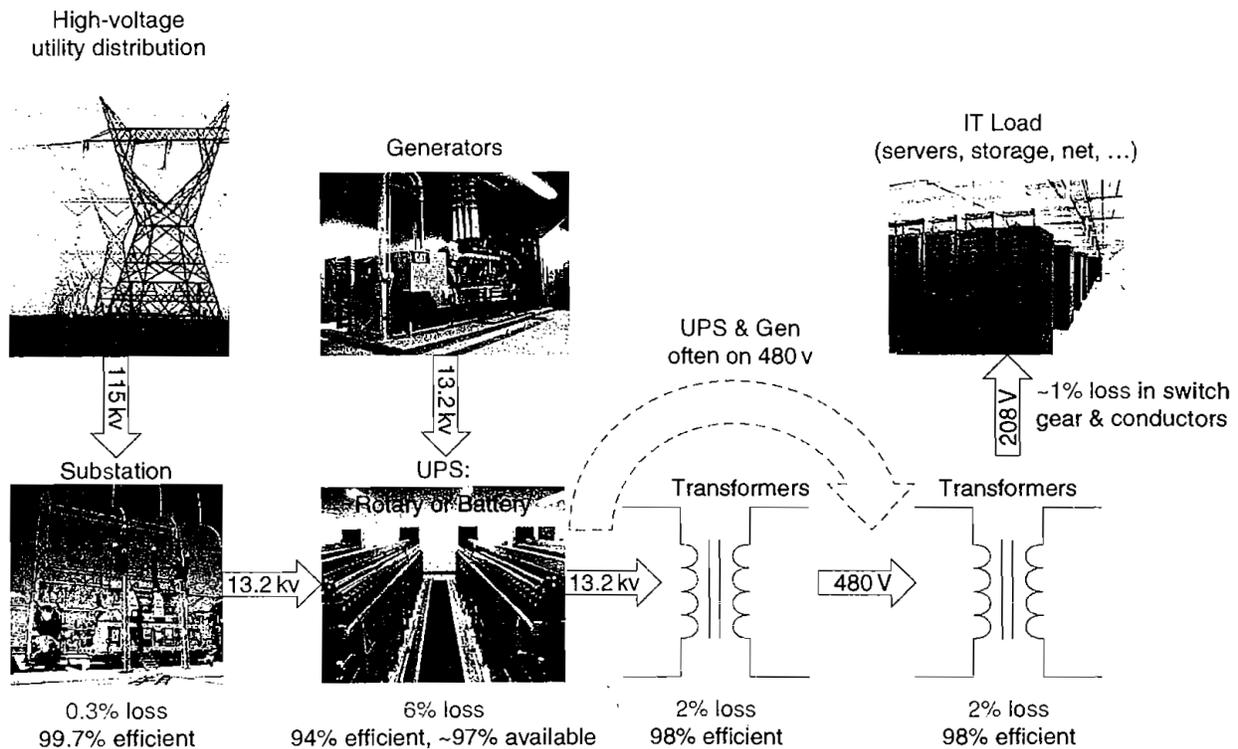


Fig: Power distribution and where losses occur

1. The substation switches from 115,000 volts to medium-voltage lines of 13,200 volts, with an efficiency of 99.7%.
2. To prevent the whole WSC from going offline if power is lost, a WSC has an uninterruptible power supply (UPS), just as some servers do. In

this case, it involves large diesel engines that can take over from the utility company in an emergency and batteries or flywheels to maintain power after the service is lost but before the diesel engines are ready. The generators and batteries can take up so much space that they are typically located in a separate room from the IT equipment. The UPS plays three roles: power conditioning (maintain proper voltage levels and other characteristics), holding the electrical load while the generators start and come on line, and holding the electrical load when switching back from the generators to the electrical utility. The efficiency of this very large UPS is 94%, so the facility loses 6% of the power by having a UPS. The WSC UPS can account for 7% to 12% of the cost of all the IT equipment.

- 3. Next in the system is a power distribution unit (PDU) that converts to low voltage, internal, three-phase power at 480 volts. The conversion efficiency is 98%. A typical PDU handles 75 to 225 kilowatts of load, or about 10 racks.
- 4. There is yet another down step to two-phase power at 208 volts that servers can use, once again at 98% efficiency.
- 5. The connectors, breakers, and electrical wiring to the server have a collective efficiency of 99%. WSCs outside North America use different conversion values, but the overall design is similar.

Putting it all together, the efficiency of turning 115,000-volt power from the utility into 208-volt power that servers can use is 89%:

$$99.7\% \times 94\% \times 98\% \times 98\% \times 99\% = 89\%$$

- The computer room air-conditioning (CRAC) unit cools the air in the server room using chilled water, similar to how a refrigerator removes heat by releasing it outside of the refrigerator. As a liquid absorbs heat, it evaporates. Conversely, when a liquid releases heat, it condenses. Air conditioners pump the liquid into coils under low pressure to evaporate and absorb heat, which is then sent to an external condenser where it is released. Thus, in a CRAC unit, fans push warm air past a set of coils filled with cold water and a pump moves the warmed water to the external chillers to be cooled down. The cool air for servers is typically between 64°F and 71°F (18°C and 22°C).
- Fig below shows the large collection of fans and water pumps that move air and water throughout the system.

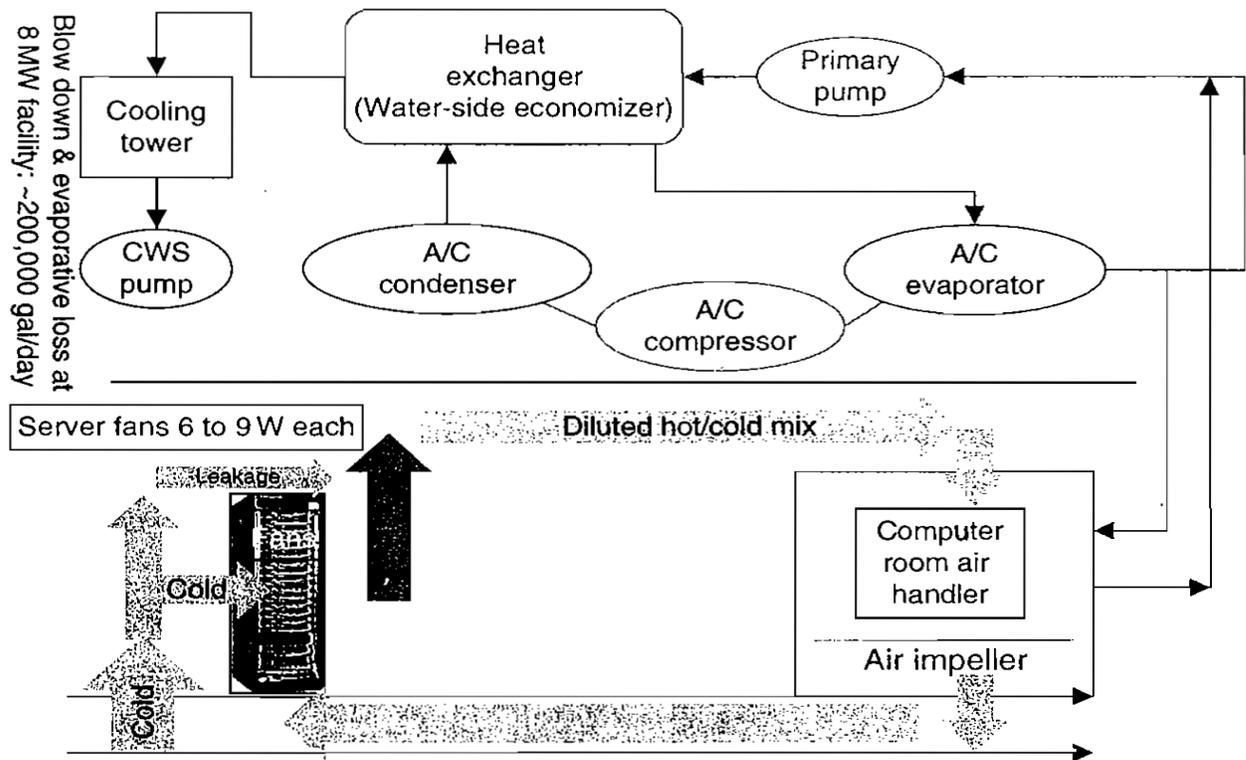


Fig: Mechanical design for cooling systems

- Clearly, one of the simplest ways to improve energy efficiency is simply to run the IT equipment at higher temperatures so that the air need not be cooled as much. Some WSCs run their equipment considerably above 71°F (22°C).
- In addition to chillers, cooling towers are used in some datacenters to leverage the colder outside air to cool the water before it is sent to the chillers. The temperature that matters is called the wet-bulb temperature. The wet-bulb temperature is measured by blowing air on the bulb end of a thermometer that has water on it.
- Warm water flows over a large surface in the tower, transferring heat to the outside air via evaporation and thereby cooling the water. This technique is called airside economization. An alternative is use cold water instead of cold air.

- Airflow is carefully planned for the IT equipment itself, with some designs even using airflow simulators. Efficient designs preserve the temperature of the cool air by reducing the chances of it mixing with hot air.
- In addition to energy losses, the cooling system also uses up a lot of water due to evaporation or to spills down sewer lines. For example, an 8 MW facility might use 70,000 to 200,000 gallons of water per day.
- The relative power costs of cooling equipment to IT equipment in a typical datacenter [Barroso and Holzle 2009] are as follows:
  - Chillers account for 30% to 50% of the IT equipment power.
  - CRAC accounts for 10% to 20% of the IT equipment power, due mostly to fans.
- Breaking down power usage inside the IT equipment itself, Barroso and Holzle [2009] reported the following for a Google WSC deployed in 2007:
  - 33% of power for processors
  - 30% for DRAM
  - 10% for disks
  - 5% for networking
  - 22% for other reasons (inside the server)

### **Measuring Efficiency of a WSC :**

- A widely used, simple metric to evaluate the efficiency of a datacenter or a WSC is called power utilization effectiveness (or PUE):
 
$$\text{PUE} = (\text{Total facility power}) / (\text{IT equipment power})$$
 Thus, PUE must be greater than or equal to 1, and the bigger the PUE the less efficient the WSC.
- Since performance per dollar is the ultimate metric, we still need to measure performance. The bandwidth drops and latency increases depending on the distance to the data. In a WSC, the DRAM bandwidth within a server is 200 times larger than within a rack, which in turn is 10 times larger than within an array. Thus, there is another kind of locality to consider in the placement of data and programs within a WSC.

- While designers of a WSC often focus on bandwidth, programmers developing applications on a WSC are also concerned with latency, since latency is visible to users. Users' satisfaction and productivity are tied to response time of a service.
- The results of experiments showed that cutting system response time 30% shaved the time of an interaction by 70%.

#### Cost of a WSC :

- WSCs costs includes operational costs as well as the cost to build the WSC. Accounting labels the former costs as operational expenditures (OPEK) and the latter costs as capital expenditures (CAPEK).
- Hamilton determined that the CAPEX of this 8 MW facility was \$88M, and that the roughly 46,000 servers and corresponding networking equipment added another \$79M to the CAPEX for the WSC.
- We can now price the total cost of energy, since U.S. accounting rules allow us to convert CAPEX into OPEX. We can just amortize CAPEX as a fixed amount each month for the effective life of the equipment.
- Note that the amortization rates differ significantly, from 10 years for the facility to 4 years for the networking equipment and 3 years for the servers. Hence, the WSC facility lasts a decade, but you need to replace the servers every 3 years and the networking equipment every 4 years.
- The fully burdened cost of a watt per year in a WSC, including the cost of amortizing the power and cooling infrastructure is :

$$\frac{\text{Monthly cost of infrastructure} + \text{monthly cost of power}}{\text{Facility size in watts}} \times 12 = \frac{\$765\text{K} + \$475\text{K}}{8\text{M}} \times 12 = \$1.86$$

- The cost is roughly \$2 per watt-year. Thus, to reduce costs by saving energy you shouldn't spend more than \$2 per watt-year. Note that more than a third of OPEX is related to power, with that category trending up while server costs are trending down over time. The networking equipment is significant at 8% of total OPEX and 19% of the server CAPEX

Size of facility (critical load watts)	8,000,000
Average power usage (%)	80%
Power usage effectiveness	1.45
Cost of power (\$/kwh)	\$0.07
% Power and cooling infrastructure (% of total facility cost)	82%
<b>CAPEX for facility (not including IT equipment)</b>	<b>\$88,000,000</b>
Number of servers	45,978
Cost/server	\$1450
<b>CAPEX for servers</b>	<b>\$66,700,000</b>
Number of rack switches	1150
Cost/rack switch	\$4800
Number of array switches	22
Cost/array switch	\$300,000
Number of layer 3 switches	2
Cost/layer 3 switch	\$500,000
Number of border routers	2
Cost/border router	\$144,800
<b>CAPEX for networking gear</b>	<b>\$12,810,000</b>
<b>Total CAPEX for WSC</b>	<b>\$167,510,000</b>
Server amortization time	3 years
Networking amortization time	4 years
Facilities amortization time	10 years
Annual cost of money	5%

Fig: Case study for a

WSC

Expense (% total)	Category	Monthly cost	Percent monthly cost
Amortized CAPEX (85%)	Servers	\$2,000,000	53%
	Networking equipment	\$290,000	8%
	Power and cooling infrastructure	\$765,000	20%
	Other infrastructure	\$170,000	4%
OPEX (15%)	Monthly power use	\$475,000	13%
	Monthly people salaries and benefits	\$85,000	2%
	Total OPEX	\$3,800,000	100%