

The Data Furnace: Heating Up with Cloud Computing

Jie Liu, Michel Goraczko, Sean James, Christian Belady
Microsoft Research
One Microsoft Way
Redmond, WA 98052

Jiakang Lu, Kamin Whitehouse
Computer Science Department
University of Virginia
Charlottesville, VA 22904

Abstract

In this paper, we argue that servers can be sent to homes and office buildings and used as a primary heat source. We call this approach the *Data Furnace* or DF. Data Furnaces have three advantages over traditional data centers: 1) a smaller carbon footprint 2) reduced total cost of ownership per server 3) closer proximity to the users. From the home owner's perspective, a DF is equivalent to a typical heating system: a metal cabinet is shipped to the home and added to the ductwork or hot water pipes. From a technical perspective, DFs create new opportunities for both lower cost and improved quality of service, if cloud computing applications can exploit the differences in the cost structure and resource profile between Data Furnaces and conventional data centers.

1 Introduction

Cloud computing is hot, literally. Electricity consumed by computers and other IT equipment has been skyrocketing in recent years, and has become a substantial part of the global energy market. In 2006, the IT industry used 61 Billion kWh electricity (or 3% of total energy consumption in the U.S.), and is the fastest growing industrial sector [9]. Energy efficiency is not only important to reduce operational costs, but is also a matter of social responsibility for the entire IT industry. The emergence of cloud computing, online services, and digital media distribution has led to more computing tasks being offloaded to service providers and increasing demand on datacenters infrastructure. For this reason, it is not surprising that data center efficiency has been one of the focuses of cloud computing and data center design and operation.

In this paper, we argue that the problem of heat generation can be turned into an advantage: computers can be placed directly into buildings to provide low latency cloud computing for its offices or residents, and the heat

that is generated can be used to heat the building. This approach improves quality of service by moving storage and computation closer to the consumer, and simultaneously improves energy efficiency and reduces costs by *reusing* the electricity and electrical infrastructure that would normally be used for space heating alone.

Physically, a computer server is a metal box that converts electricity into heat¹. The temperature of the exhaust air (usually around 40-50°C) is too low to regenerate electricity efficiently, but is perfect for heating purposes, including home/building space heating, cloth dryers, water heaters, and agriculture. We propose to replace electric resistive heating elements with silicon heating elements, thereby reducing societal energy footprint by using electricity for heating to also perform computation. The energy budget allocated for heating would provide an ample energy supply for computing. For example, home heating alone constitutes about 6% of the U.S. energy usage². By piggy-backing on only half of this energy, the IT industry could double in size without increasing its carbon footprint or its load on the power grid and generation systems.

Technological and economical trends also make energy reuse a promising direction. After years of development of cloud computing infrastructure, system management capabilities are getting mature. Servers can be remotely re-imaged, re-purposed, and rebooted. Virtual machine encapsulation ensures certain degree of isolation. Secure executions on untrusted devices are feasible. Sensor networks have made high physical security within reach. At the same time, computers are getting cheaper, network connectivity is getting faster, yet energy becomes a scarce resource and its price is on the

¹At the same time, some energy is used by the fans to move the air for cooling and a very small amount of energy is used to drive electrons or photons on the network interface

²According to the Energy Information Administration (EIA), U.S. households used 154 billion kWh for home heating, 116 billion kWh of which goes to the heating element[4]

trend of fast rise.

From a manageability and physical security point of view, the easiest adopters of this idea are office buildings and apartment complexes. A mid-sized data center (e.g. hundreds of kilowatts) can be hosted inside the building and the heat it generates will be circulated to heat the building. Dedicated networking and physical security infrastructure can be built around it, and a dedicated operator can be hired to manage one or more of them. Their operation cost will be similar to operating other urban data centers, and can leverage the current trend toward sealed server containers that are replaced as a unit to save repair/replacement costs.

As a thought provoking exercise, we push this vision to the extreme in this paper. We investigate the feasibility of *Data Furnaces* or DFs, which are micro-datacenters, on the order of 40 to 400 CPUs, that serve as the primary heat source for a single-family home. These micro-datacenters use the home broadband network to connect to the cloud, and can be used to host customer virtual machines or dedicated Internet services. They are integrated into the home heating system the same way as a conventional electrical furnace, using the same power system, ductwork, and circulation fan. Thus, DFs reduce the cost per server in comparison to conventional data centers by leveraging the home’s existing infrastructure, and precluding the cost of real estate and construction of new brick and mortar structures. Furthermore, they naturally co-locate computational power and storage close to the user population.

DFs are managed remotely and do not require more physical maintenance than conventional furnaces. The cloud service operators may further incentivize the host families by providing free heat in exchange for occasional physical touches such as replacing air filters and, in extreme cases, turning on/off servers.

Previous studies also suggest bringing micro-datacenters close to the users, including the idea of renting condos to host servers [3] and to use home routers as a nano-datacenter for content caching [10]. In this paper, we suggest that a quantum leap is achieved when this idea scaled to the size of a furnace; at this scale, the micro-datacenter cannot only leverage existing residential infrastructure for power, networking, and air circulation, but it can also reuse the energy that would otherwise be consumed for home heating.

In the rest of the paper, we first compare the cost of Data Furnaces versus building and operating conventional data centers, by analyzing the heating demand in various climate zones to understand the capacity and utilization requirements. Then, in Section 3 we discuss what kind of services can run on DFs given different investment levels. Finally, we discuss technical challenges including security, power management and system man-

Climate Zones	Locations	Outdoor Temperature	
		< 21°C	> 35°C
Zone 1	Minneapolis, MN	82.20%	0.11%
Zone 2	Pittsburgh, PA	82.46%	0.00%
Zone 3	Washington, D.C.	77.22%	0.13%
Zone 4	San Francisco, CA	96.44%	0.00%
Zone 5	Houston, TX	46.51%	0.15%

Table 1: Representative locations used in simulations




Figure 1: Average monthly heating rate across the year

agement.

2 Home Heating with Data Furnaces

In this section, we evaluate the financial viability of Data Furnaces from the perspective of cloud service providers. Because DFs serve as a primary heat source in homes, we first perform a simulation study to understand the heating demands for a single family house across the climatic zones in the U.S.. Based on the results, we discuss the expected savings if DFs were used in each zone. We use ballpark figures and back-of-the-envelope calculations; the exact numbers depend on the specific households and data centers under consideration.

DFs reduce the total cost of conventional datacenters in three main ways. First, much of the initial capital investment to build the infrastructure for a datacenter is avoided, including real estate, construction costs, and the cost of new power distribution, networking equipment, and other facilities. A second and related benefit is that operating costs are reduced. For example, cooling cost is significant in centralized data centers due to the power density [13], but DFs have essentially no additional cooling or air circulation costs since the heat distribution system in the house already circulates air. Thus, DFs increase the power usage efficiencies (PUE) over conventional datacenters. Finally, the money to buy and operate

a furnace for home heating is avoided, and can be used instead to offset the cost of servers: the cloud service provider can sell DFs at the price of a furnace, and charge household owners for home heating. By doing this, the heating cost remains the same for the host family, while costs are reduced for the cloud service provider.

One disadvantage of DFs is that the retail price of electricity is usually higher in the residential areas by 10% to 50% than industrial areas [8]. Another potential disadvantage is that the network bandwidth can cost more in homes if the home broadband link cannot satisfy the service and a high bandwidth link must be purchased. Finally, maintenance costs will increase because the machines will be geographically distributed.

To weigh these advantages and disadvantages, we perform a total cost of ownership (TCO) analysis for both DFs and conventional data centers. The initial and operating cost can change based on climate zone, so we first measure the actual heating demand for homes using the U.S. Department of Energy’s EnergyPlus simulator [2]. This simulator calculates the heating load (BTU) required each minute to keep a home warm, using actual weather traces recorded at airports. We simulate a 1700 square foot residential house that is moderately insulated and sealed with a heating setpoint of 21°C (70°F). We use weather data of Typical Meteorological Year 3 (TMY3) [12], and replay the entire year for cities in each of the five climate zones in the U.S., as listed in Table 1. The last two columns show the percentage of time (in minutes granularity) that the outside temperature is less than 21°C, (thus heating is useful) and that the outside temperature is greater than 35°C (thus the server may have to be shut down for thermo protection since we do not expect cooling the furnace.). The percentage of time in between is when the servers can be run but the heat must be pumped outside.

Table 2 summarizes the results of our TCO analysis for a DF in each climate zone. Row 1 indicates the maximum heating rate required for the coldest day of the year in each zone, based on our simulations. Row 2 indicates the number of servers required in a DF to satisfy these requirements, assuming a typical server such as the Dell PowerEdge 1850 consumes 350 W generating 1,194.55 BTU/h heat dissipation [1]. The third row shows the estimated home heating cost for each zone, assuming the consumer pays an amount equal to the cost of operating either an appropriately sized oil furnace in Zones 1 and 2, or a heat pump in Zones 3, 4 and 5. This number is calculated based on the data in Figure 1, which shows the average heating rate required for each month in each climate zone, as determined by our simulations. The actual home heating cost certainly varies based on the heating technology and thermostat setting. This is the incentive to the host family that the cloud operator can claim.

The fourth row shows the income per year from selling the DF to the consumer, assuming the consumer pays \$4000 for a comparably sized heat pump or furnace. Based on the electricity price data, the fifth row shows the increased cost of electricity for the data center owner, due to residential electricity rates being higher than commercial electricity rates; we assume that the residential electricity rate in the U.S. is \$0.10/kWh while the industrial rate is \$0.05/kWh. The sixth and seventh rows indicate the cost per DF for each zone, and the cost per server based on the number of servers in each zone. These numbers represent overhead costs and do not include the costs of the servers themselves, since those costs must also be paid in a conventional data center. The eighth row indicates the optional cost of upgrading the network link at the home to a T1, which is assumed to cost \$2,640 for the home per year. In all regions, the overhead cost is less than \$90 per server per year.

In comparison, the overhead cost per server per year for a conventional datacenter is approximately \$400. This cost is taken from James Hamilton’s spreadsheet analysis [5], which estimates the cost of running a datacenter of over 45,000 servers, using a 3-year server amortization and a 15-year facilities amortization. These overhead costs include construction for a brick and mortar building, network equipment, cooling costs, and other equipment and operational costs that are not needed for DFs. Based on these numbers, the last row of Table 2 indicates the cost savings per server per year of running a DF, in comparison to a conventional datacenter. These savings would be higher for data centers that run applications that do not necessitate an upgrade to the network connection, or when comparing to urban data centers that pay more for electricity and real estate.

Designed correctly, Data Furnaces may not require frequent touch for maintenance. Studies have shown that on average 9% servers fail in 14-month of operation [11]. In other industries, the cost to roll a truck is approximately \$100 per maintenance visit, which is only necessary for a Data Furnace when a few physical failures occur that cannot be managed remotely. Even if these occur three times per year, it is an amortized cost of less than ten dollars per server per year and therefore does not substantially change the TCO analysis.

3 Types of Service

The overhead costs for Data Furnaces are very small compared to centralized mega datacenters, but the true value proposition depends on the types of services it can provide. DFs will present both opportunities and challenges for cloud computing services. On one hand, DFs can improve performance for many applications because they are close to areas of high population density. On

		Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
Heating Requirements	Sizing for coldest day of year (BTU/h)	133,790	136,179	120,650	54,949	44,198
	Number of 350W Servers in Furnace	112	114	101	46	37
	Est. normal home heating cost	\$2956.9	\$2090.3	\$2450.0	\$1499.5	\$711.0
Amortized cost/gain (per furnace per year)	Income from selling DF	-\$400.0	-\$400.0	-\$400.0	-\$400.0	-\$400.0
	Cost increase due to residential rates	\$9525.4	\$6733.6	\$5742.7	\$3514.9	\$1666.5
Amortized overhead expense per server per year		\$81.5	\$55.56	\$52.9	\$67.7	\$34.2
Optional Cost (per year)	Amortized network upgrade per server	\$22.0	\$21.6	\$24.4	\$53.5	\$66.5
Savings per server per year (compare to \$401.6 per server for conventional DC)		\$298.1	\$324.4	\$324.3	\$280.4	\$300.9

Table 2: Cost benefit to run a DF over a conventional datacenter (from the perspective of cloud service provider)

the other hand, the new cost structure for bandwidth and energy may make DFs less attractive for other types of applications. We envision at least three classes of DFs, as described below:

Type A: Low-cost seasonal DFs

The least expensive DFs will use the existing home broadband connection and will perform computation predominantly when the house needs heating (e.g. at nights or during winters), providing seasonal computing capacity with almost no operational cost. These low-cost DFs can also reduce initial costs by using recycled, older models of servers – hundreds of thousands of which are decommissioned from existing datacenters every year to make space for servers with latest hardware technology. These older servers are often less energy efficient and less cost effective in a conventional datacenter, but ideal for providing heat and maintaining basic disk and network operations year-round. Because of the population distribution on the planet, the computing capacity of this class of DFs will be skewed toward the northern hemisphere’s winter months. Despite their bandwidth-limited and seasonal nature, Type A DFs can still support a wide range of services. For example, many delay-tolerant batch jobs can be performed opportunistically, such as non-real-time web crawling, content indexing, and the processing of large scientific data sets (e.g. astronomical data, genome sequencing operations, and SETI@Home). Type A DFs can also help developing communities, societal services, hobbyists, and other organizations looking for extremely low cost computing resources. Finally, governments currently subsidize home heating for developing communities³. With the same fund, Type A DFs can have extended operation time for education, scientific, and social services, doubling the benefits.

Type B: Low-bandwidth neighborhood DFs

A lower-cost DF can also operate year-round using

³e.g. Wisconsin Home Energy Assistance Program (WHEAP) (<http://homeenergyplus.wi.gov/category.asp?linkcatid=239&linkid=118>)

the broadband channel of the home instead of upgrading the network link, thereby supporting typical computation loads but with a relatively slow egress link. The least expensive cable modem package from Comcast today provides 12 Mbps download and 2 Mbps upload speed, in comparison to a conventional data center, which may have a 10 Gbps link (shared among hundreds of thousands of machines). This new resource profile makes financial sense for DFs, particularly in warmer climates, where it will reduce overhead costs by up to 80% per server, as shown in Table 2. Despite bandwidth limitations, Type B DF can nevertheless provide high quality computing services due to geographic proximity to the customers. For example, Internet television services and on-line video rental services could use pre-fetching based on local programming schedules or video queues of people in the local vicinity. Similarly, location-based services such as map serving, traffic estimation, local navigation, and advertisements for local stores are typically requested by customers from the local region. Network latency can be reduced by creating more data centers closer to customers [7], particularly for cloud applications such as E-mail, multimedia distribution [3, 10], and low-hit rate Web pages and services that are delay tolerant, embarrassingly distributable, and/or can benefit from large-scale replication and geo-distribution.

Type C: Eco-friendly urban DFs

Closest to the current operation modes, high-end DFs will upgrade the communication link at the home to FiOS or a T1, and can operate year-round even when heat is not needed in the home by pumping excess heat directly outside. This is particularly promising for business buildings and apartments. Having comparable resources per server, these DFs would give service providers the ability to expand into urban areas more quickly and easily without urban real estate and infrastructure expenses, as long as the application scale to the number of servers. On the other hand, Type C DFs also present new challenges. First, summer operation costs more because higher elec-

tricity rates are not compensated by sale of heat. However, this additional cost would be marginal compared to the cost advantages of DFs, especially since many data centers only have about 20% server utilization anyway. Furthermore, urban areas present opportunities for larger DFs in commercial buildings, business parks, and apartment complexes where lower commercial energy rates are available and where the larger heating units have an energy footprint close to containerized datacenters that are used for edge services. A second challenge is that DFs may break large services into many small pieces, thereby introducing a different granularity to perform massively distributed computations, like MapReduce.

4 Technical Challenges

Despite the facts that server hardware is getting more robust and less sensitive to operation conditions, and that cluster-level management software is getting mature, DF present new unique challenges due to its distribution in residential households, including:

Isolation: DFs use the existing power and network infrastructure of a normal residence or business, and they must not affect normal building operations. A DF must monitor and react to the total power and broadband usage of its hosts in order not to exceed capacity (e.g. 24kW and 12Mbps for a typical home in Seattle), and cloud services must be designed to accommodate this low-priority operation.

Security: One difference between DFs and dedicated datacenters is the challenge of physical security, and DFs should be treated as being in the most insecure environment. For example, each server should have an individual tamper proofing device, such as a networked sensor. This will allow them to be individually swapped out by the hosting party only when necessary (e.g. due to unrecoverable device failure). Furthermore, all stored data and network traffic must be encrypted. Software running on the servers should be sandboxed and secured from the hosting party.

Zero-touch management: Automated solutions exist for monitoring server and network health, remotely re-purpose and re-image machines, and remotely diagnose faults and power cycle machines (e.g. [6]). However, without dedicated system operators and in a residential environment, DFs require high heating reliability and zero-touch management. Even at the event of software failure, the system should continue to provide heat until receiving physical services.

5 Conclusion

With continuously increasing demand for computing, the community must explore novel ways to expand computing capacity without increasing financial burden and energy costs. Data Furnaces will reuse the facilities and energy already allocated for heating purposes to provide computing services with low cost and energy footprint. In this paper, we focus on homes as an illustrating example, but a similar approach could be used to heat water tanks, office buildings, apartment complexes, vegetable farms, and large campuses with central facilities.

References

- [1] Dell PowerEdge 1850 Environmental Data Sheet. http://www.dell.com/downloads/global/corporate/environ/PE_1850.pdf.
- [2] EnergyPlus Energy Simulation Software. <http://apps1.eere.energy.gov/buildings/energyplus/>.
- [3] CHURCH, K., GREENBERG, A., AND HAMILTON, J. On delivering embarrassingly distributed cloud services. In *ACM SIGCOMM Hotnets VII, Calgary Alberta, Canada* (2008).
- [4] ENERGY INFORMATION ADMINISTRATION, IN THE US DEPARTMENT OF ENERGY. End-Use Consumption of Electricity Survey. <http://www.eia.doe.gov/emeu/recs/recs2001/enduse2001/enduse2001.html>, 2001.
- [5] HAMILTON, J. Overall Data Center Costs. <http://perspectives.mvdirona.com/2010/09/18/OverallDataCenterCosts.aspx>.
- [6] ISARD, M. Autopilot: automatic data center management. *SIGOPS Oper. Syst. Rev.* 41 (April 2007), 60–67.
- [7] LI, A., YANG, X., KANDULA, S., AND ZHANG, M. Cloudcmp: Comparing public cloud providers. In *Internet Measurement Conference, Melbourne, Australia* (2010).
- [8] U.S. ENERGY INFORMATION ADMINISTRATION. Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State.
- [9] U.S. ENVIRONMENTAL PROTECTION AGENCY ENERGY STAR PROGRAM. Report to congress on server and data center energy efficiency public law 109-431. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.
- [10] VALANCIUS, V., LAOUTARIS, N., MASSOULIE, L., DIOT, C., AND RODRIGUEZ, P. Greening the internet with nano data centers. In *ACM CoNEXT'09, Rome, Italy* (2009).
- [11] VISHWANATH, K. V., AND NAGAPPAN, N. Characterizing cloud computing hardware reliability. In *Proceedings of the 1st ACM symposium on Cloud computing* (New York, NY, USA, 2010), SoCC '10, ACM, pp. 193–204.
- [12] WILCOX, S., MARION, W., AND NATIONAL RENEWABLE ENERGY LABORATORY (U.S.). *Users Manual for TMY3 Data Sets*. National Renewable Energy Laboratory, 2008.
- [13] WOODS, A. Cooling the data center. *ACM Queue* 8, 10:10–10:19.