



Google's Green Computing: Efficiency at Scale

Introduction

Powering an Email System

Provisioning Email Servers
Redundancy and Reliability
Computer Housing and Cooling

Cloud Advantages

Carbon Costs

Conclusions

Introduction

It's common to hear about new data centers being built, and it may seem as if the energy used by "the cloud" is a growing problem. However, services provided by the cloud can be remarkably efficient. In many cases, data centers hosting cloud services are more efficient than the in-house office servers they replace. These efficiency gains come from maximizing server utilization, using high-efficiency facilities and making power usage efficiency a priority for hardware and software developers.

This case study compares the energy savings and carbon footprint of using Gmail via Google Apps—Google's cloud-based messaging and collaboration suite, versus housing local servers to manage the same email. **Google's cloud services are all carbon neutral**—dating back to 2007. However, for the purposes of this case study, we will be modeling Google's per-user carbon footprint before offsets are applied.

Powering an Email System

When a user checks their email, energy is consumed in three places:

1. **The client:** The PC, laptop, phone or other device directly accessing email is called the client. It's using electricity as it fetches and displays email.
2. **The network:** Wireless routers, network switches and all the intervening networking equipment between the client and the email server use electricity. In a small office, a wireless router may be the only network device involved. In a larger enterprise, networking equipment could involve dozens of devices spanning continents.
3. **The server:** The computer or group of computers receiving, sending and storing email is called the email server. This machine or group of machines is constantly on and consuming electricity.

Although switching from a locally hosted email system to a cloud-based email system can affect the energy usage of all three, it primarily affects the server level energy usage.¹ This paper examines server energy only. Servers in locally hosted email systems are often underutilized and installed in facilities that aren't optimized for energy efficiency. Servers in cloud-based email systems are located in large, efficient data centers which can provision more servers as needed. A cloud-based email system saves considerable amounts of per-user energy costs once provisioning email servers, providing redundancy, and cooling costs are taken into account.

Provisioning Email Servers

An email server is a computer that sends outgoing messages to the external network and stores incoming messages until a user downloads or deletes them. Servers can be very large, powerful machines for handling many simultaneous users or smaller machines for handling only a few.

Consider three typical businesses that choose to host their mail locally and their hypothetical server requirements²:

Business Size	Email Users	Server Requirements
Small	50	A single, mid-range multi-core server with local disk that can serve 300 ³ users and draws 200 ⁴ Watts.
Medium	500	A single, large, many core server with combinations of local and network storage, which can host 1,000 users and which draws 450 Watts. ⁵
Large	10,000	Several, large, many-core servers with combinations of local and network storage which can host 1,000 users and draws 450 Watts.

The small organization is at a distinct disadvantage, having to purchase a server that can host 300 users when it only needs capacity for 50. An organization can't purchase and power half a server if that's all they need. Moreover, the energy consumption of a fully-utilized server is not much higher than a partly-utilized one.

Redundancy and Reliability

A modern day business relies on email servers 24 hours a day, 7 days a week, 365 days a year and endeavors to make them as reliable as possible. Depending on the size and type of the organization, this could mean:

- Additional backup email servers
- Redundant email storage, like networked hard drives or other storage technology
- Backup networking links
- Co-locating email servers in multiple separate physical locations

"N" shorthand is sometimes used to describe the number of redundant servers needed for an organization's email. If "N" number of servers are need to handle an organization's users, "N+1" redundancy means there is one spare server on hand. If a company requires two servers to host all of their email, then they have a third on hand in case one fails. "2N" redundancy would mean having one spare for every required server.

Redundancy means more reliability and more energy use. Modeling modest redundancy requirements for a small, medium and large business provides an estimated number of servers each business will need for their email system.

Business type	Minimum servers required without redundancy	Redundancy plan	Total servers required
Small	1	N+1	2
Medium	1	N+1	2
Large	10	N+0.2N	12

From the number of servers, the annual energy usage can be roughly calculated as such:

Business type	Individual server power	Total server power	Server power per user	Annual server energy required per user
Small	200 W	400 W	8 W	70 kWh
Medium	450 W	900 W	1.8 W	16 kWh
Large	450 W	5400 W	0.54 W	4.7 kWh

The disadvantages faced by small and medium businesses are magnified once redundancy requirements are added. The larger the organization, the more efficiently servers can be provisioned. When comparing server power per user, the large organization uses 1/15th as much energy per user as the small organization.

Computer Housing and Cooling

The energy email servers consume directly is only one component of the energy used to store and deliver email. Additional energy is required to house and cool the servers. A server located under a desk or in an office networking closet makes use of the same air conditioning system as people in the office. This can be very inefficient, since the room temperature and ventilation are set up for personal comfort, not for server efficiency. It is not uncommon for such a system to use an additional 1.5 W of extra energy overhead for every 1 W of computer load.

The data center industry uses the measurement “PUE,” or power usage effectiveness, to calculate the energy costs of housing and cooling servers. PUE measures how much overhead energy is required to house and cool the computers inside a building relative to the amount the computers consume themselves⁶. For the example above, the email server in the networking closet has a PUE of 2.5. That’s 2.5 watts of power total for each watt of computing power.

Medium and large data centers are typically a bit better—and can be much better. This is because they are designed with powering and cooling IT equipment in mind rather than keeping humans comfortable. The average PUE of a typical data center today is about 1.8⁷, but for a larger organization with time to make significant measurements and improvements, it can be lower. For our example, we estimated 1.6.

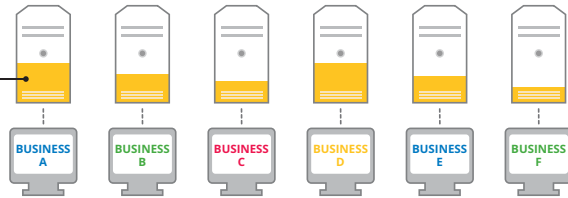
Business Type	IT power per user	PUE	Total power per user	Annual energy per user
Small	8 W	2.5	20 W	175 kWh
Medium	1.8 W	1.8	3.2 W	28.4 kWh
Large	0.54 W	1.6	0.9 W	7.6 kWh

Once again, small setups underperform, with both larger per-user energy costs and carbon footprints. When cooling and housing costs are taken into account, the total power per user for a large organization can be 1/20th that of a small business.

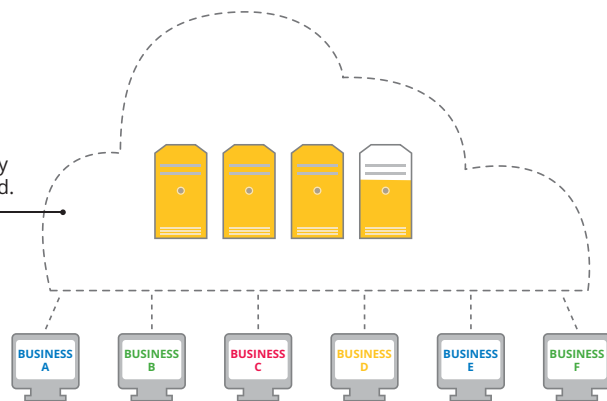
Cloud Advantages

So far, our model shows that large organizations hold an advantage over smaller organizations: they can provision resources more efficiently and they can operate those computers in more efficient facilities. Cloud providers take advantage of this efficiency in scale by providing servers for millions of users—maximizing the utilization of machines while cutting down on the total number of servers required. The result is fewer machines and less energy over all.

Locally Hosted Email:
Inefficient use of servers by individual businesses.



Cloud-Based Email:
Efficient use of servers by collective use in the cloud.



For example, Google has developed Gmail for Google Apps, a cloud-based email service used by over 4 million organizations, small and large. Gmail is constantly improving as we integrate new efficiency improvements, develop ways of streamlining our operations, and, of course, increase capabilities. Software developers, hardware designers and data center technicians have a unified goal of optimizing Google services to use as few resources as possible.

This is quite distinct from the usual software and server model, where software developers from one company have to develop an email server that will work on wide range of hardware developed by other companies. At Google, we can optimize across boundaries that are not accessible in a traditional software or hardware development environment.

Our size, focus, and ability to optimize across barriers translates into a series of distinct advantages for our servers hosting cloud-based email, including:

- **Custom, high-efficiency servers:** Our servers are specifically designed to host cloud-based services with only the components that are necessary for them to perform—nothing more or less.
- **Custom, high-efficiency power supplies:** Many servers today use power supplies that have efficiencies in the mid-80% range. That means for every 100 Watts the computer needs, almost 20 Watts are converted directly to heat. Google's are more than 90% efficient, wasting less energy.
- **Custom-built software:** From top to bottom, Google infrastructure is designed to deliver our innovative web services. Our software is designed to perform with maximum efficiency on the servers we develop without bloated additional features.

Optimizing the entire process of storing, hosting and serving email means that Gmail requires less than 250 mW per user⁸. Over the course of a year, that's just a bit more than 2 kWh of energy, or about \$0.22 per user per year,⁹ representing significant savings over the locally hosted alternatives modeled in the previous example.

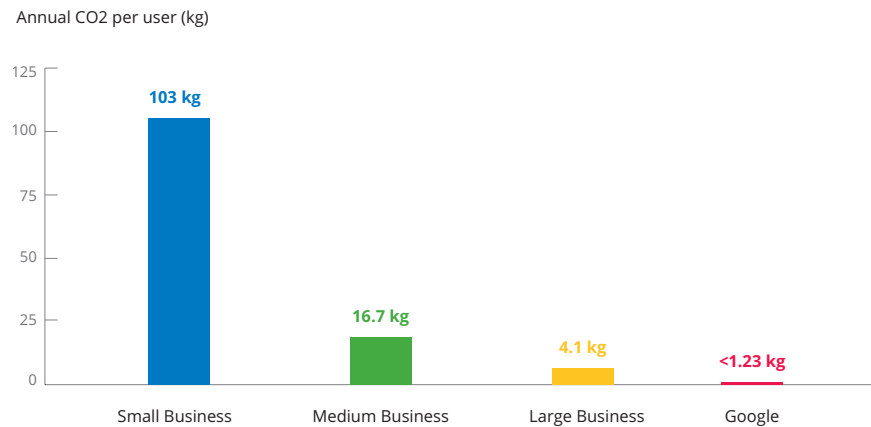
Business Type	IT power per user	PUE	Total power per user	Annual energy per user
Small	8 W	2.5	20 W	175 kWh
Medium	1.8 W	1.8	3.2 W	28.4 kWh
Large	0.54 W	1.6	0.9 W	7.6 kWh
Gmail	< 0.22 W	1.16	< 0.25 W	< 2.2 kWh

Carbon Costs

Since 2007, Google has been a carbon neutral company, utilizing energy efficient improvements, green power and carbon offsets to bring our footprint down to zero. Users of Gmail also enjoy this benefit of zero carbon emission email.

However, to form an accurate comparison of cloud-based versus locally hosted services, we calculated our per-user carbon emissions before offsets. Using the **EPA** US average emissions intensity for locally hosted email servers¹⁰, and using Google's own calculated emissions intensity for its operations, the annual carbon footprint of a Gmail user is about 1/80th that of a small business with locally hosted email servers.

Locally Hosted Server vs. Gmail: A comparison of annual CO2 per user



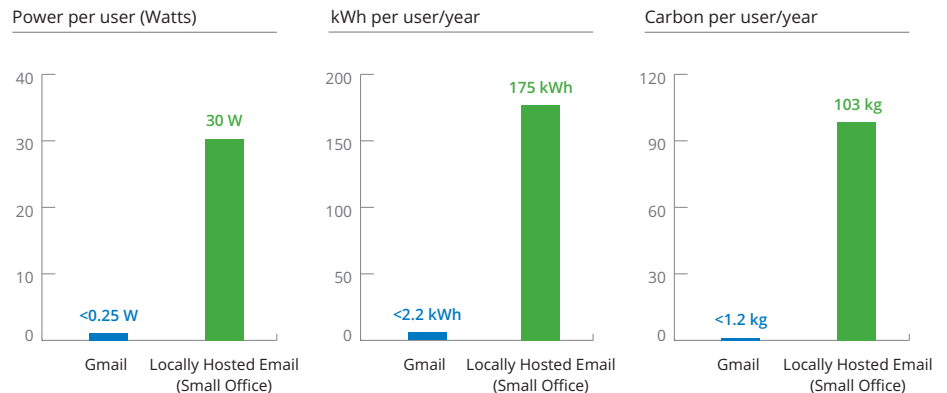
* Small Business = 50 email users, Medium-sized = 500 email users, and Large = 10,000+ email users

Conclusions

When hosting services, larger organizations have the benefit of scale. However, there are many more small businesses than large ones. In the United States more than 18% of employees work for businesses with fewer than 20 employees and more than 50% work for businesses with fewer than 500.¹¹

Cloud-based services like Gmail allow organizations of all sizes to reap these scale advantages of increased efficiency, reduced overhead costs, and smaller carbon footprint without needing the expertise of an army of software developers, hardware designers and data center technicians. For a small office of 50 people, choosing Gmail over a locally hosted server can mean an annual per-user power savings of up to 170 kWh and a carbon footprint reduction of up to 100 kg of CO₂.¹² Larger organizations show smaller, though still impressive efficiency gains.

Locally Hosted Email (small office*) vs. Gmail: Comparison of power, energy and carbon use per user/year



* Small office = 50 users

Gmail has shown for years that the cloud can deliver a high-quality, reliable, and useful service at a much lower energy cost than other methods. File storage, calendar, teleconferencing, voicemail, chat and document management all enjoy these energy economies. Even traditional applications such as word processing and spreadsheets may ultimately benefit if users transition from traditional PCs and laptops to lower energy devices like tablets or netbooks that store their information in the cloud.

Cloud-based services are growing. Email, movies, music, television and telephone services increasingly rely on cloud computing to serve, store and transmit data. As these technologies develop and mature, efficient hardware, software, and server provisioning will continue to make the cloud the most energy efficient platform for delivering computing.

- 1 There is no appreciable difference in client energy, since the user is usually not changing the device they use to access email. We would expect network energy to increase somewhat, as more traffic must traverse the Internet in the cloud-based solution. However, this effect is secondary to the large effect on server energy.
- 2 These are rough guesses at machine configurations. The actual provisioning of an enterprise email system is exceedingly complex, with many tradeoffs to be made. For example, Microsoft **provides a calculator** to assist administrators in appropriately provisioning Exchange servers.
- 3 A single, modestly equipped server can manage this load, depending on user behavior and mailbox size. For example, a server with four 1 TB disk drives configured with **RAID 5** redundancy would have 3 TB available for storage, enough to store about 600 5 GB email accounts.
- 4 Estimate is of average power draw of a typical 2U server configured with 4 TB of disk, 64GB memory, and dual Xeon processors. Servers vary, but exploration the **Dell Energy Smart Solution Advisor** and the **HP Power Advisor**, both online tools of estimating server power, indicate that 200-400 W is a reasonable range.
- 5 450 W was chosen because it represents the total power of all servers and storage in the world, averaged on a per-server basis. Data is from Koomey 2011 **Growth in data center electricity use 2005 to 2010**. Oakland, CA: Analytics Press. July 2011.
- 6 More information on PUE, including how we apply it in our own data centers, is available in the **Google's Green Data Centers: Network POP Case Study** white paper.
- 7 From the **Uptime Institute survey of 500 data centers**.
- 8 This is using a detailed estimate of all the energy necessary to host Gmail divided by our 7-day active users as our user count. There are many more Gmail accounts than 7-day active users, but this best represents accounts in regular use.
- 9 Using \$0.1017/kWh, the **EPA's US average commercial rate for 2009**.
- 10 Emissions factors vary significantly geographically, so a given company's emissions could be higher or lower depending on where they are located.
- 11 2005 data from **The Small Business Economy, Report to the President 2008**, Appendix A.
- 12 Google has been a carbon neutral company since 2007. We calculate the amount of carbon released by our electricity usage and "zero" out our carbon footprint using renewable energy purchases and carbon offsets. This is the number we calculate which is then zeroed out by offset purchases.

