

**Computing****Down on the server farm**

May 22nd 2008 | SEATTLE  
From The Economist print edition

**The real-world implications of the rise of internet computing**

EVEN when the sky is blue over Quincy, clouds hang in the air. The small town in the centre of the state of Washington is home to half a dozen huge warehouses that power the global “computing clouds” run by internet companies such as Yahoo! and Microsoft. The size of several football pitches, these data centres are filled with thousands of powerful computers and storage devices and are hooked up to the internet via fast fibre-optic links.

Yet even more intriguing than the buildings' size is their location. Quincy is literally in the middle of nowhere, three hours' drive from the nearest big city, Seattle. But it turns out to be a perfect location for data centres. As computing becomes a utility, with services that can be consumed from everywhere and on any device, ever more thought is being put into where to put the infrastructure it needs.

Where the cloud touches down is not just the business of the geeks. Data centres are essential to nearly every industry and have become as vital to the functioning of society as power stations are. Lately, centres have been springing up in unexpected places: in old missile bunkers, in former shopping malls—even in Iceland. America alone has more than 7,000 data centres, according to IDC, a market-research firm. And each is housing ever more servers, the powerful computers that crunch and dish up data. In America the number of servers is expected to grow to 15.8m by 2010—three times as many as a decade earlier.

Until a few years ago, the location of servers was an afterthought, says Jonathan Koomey, a consulting professor of environmental engineering at Stanford University. Most sat in cupboards or under desks. The computers in corporate data centres were often housed in the firm's basement. And dedicated “server farms”, which came of age during the dotcom bubble and often housed the machines of internet start-ups, were mostly built in Silicon Valley and other high-tech hubs.

**The geography of the cloud**

Now this haphazard landscape is becoming more centralised. Companies have been packing ever more machines into data centres, both to increase their computing capacity and to comply with new data-retention rules. As space gets tight and energy costs climb, many firms have begun consolidating and simplifying their computing infrastructure. Hewlett-Packard, the world's biggest computer-maker, for instance, is replacing its 85 data centres across the world with just six in America.

Internet firms, meanwhile, need ever larger amounts of computing power. Google is said to operate a global network of about three dozen data centres with, according to some estimates, more than 1m servers. To catch up, Microsoft is investing billions of dollars and adding up to 20,000 servers a month.

As servers become more numerous, powerful and densely packed, more energy is needed to keep the data centres at room temperature. Often just as much power is needed for cooling as for computing. The largest data centres now rival aluminium smelters in the energy they consume. Microsoft's \$500m new facility near Chicago, for instance, will need three electrical substations with a total capacity of 198 megawatts. As a result, finding a site for a large data centre is now, above all, about securing a cheap and reliable source of power, says Rich Miller of Data Center Knowledge, a website that chronicles the boom in data-centre construction.

The availability of cheap power is mainly why there are so many data centres in Quincy. It is close to the Columbia River, with dams that produce plenty of cheap hydroelectric power. There is water for cooling, fast fibre-optic links, and the remoteness provides security. For similar reasons, Google chose to build a new data centre at The Dalles, a hamlet across the Columbia River in the state of Oregon.

Such sites are in short supply in America, however. And with demand for computing picking up in other parts of the world, the boom in data-centre construction is spreading to unexpected places. Microsoft is looking for a site in Siberia where its data can chill. Iceland has begun to market itself as a prime location for data centres, again for the cool climate, but also because of its abundant geothermal energy. Hitachi Data Systems and Data Islandia, a local company, are planning to build a huge data-storage facility (pictured at top of article). It will be underground, for security and to protect the natural landscape.

So will all data centres end up in remote places like Quincy or Iceland? Not necessarily. For many applications, speed is of the essence. To make sure that its web-search results show up almost instantly, Google has to distribute its data centres around the world. Financial-services firms want to have access to trading data in real time, which explains the high density of data centres near New York and London. And fast-moving online games must be hosted near their players.

Even when speed does not matter, many firms want their servers close by, says Mike Foust, the boss of Digital Realty Trust, which builds and rents out data centres. Sometimes this is for maintenance; sometimes it is because "server huggers" do not want to let go. Security also counts. The Boyd Company, which advises companies about where to put their data centres, thinks more should be built in the provinces. Demand for secure locations for back-up centres, which many firms now have to maintain, will give rise to huge regional data centres, such as the one being built in Newport in Wales.

The criteria that companies use to pick a site keep evolving, says Mike Manos, Microsoft's director of data centres. His team feeds 35 sets of data into an electronic "heat" map of the world. With colours that range from green to red, it shows where conditions are favourable and which places should be avoided. And Microsoft needs a lot of choice: if a new service suddenly becomes popular, it needs to be able to increase computing capacity quickly. That is also why it uses shipping containers pre-loaded with up to 2,000 servers, which can be up and running within hours. In the firm's Chicago data centre, over 200 such containers will populate an entire floor.

Yet it will not just be market economics that determines the shape of the clouds. Local governments give tax breaks in the hope that the presence of big data centres will attract other businesses (the computing plants themselves usually employ only a few dozen people). Regulation is a factor, too. SWIFT, a bank-transfer consortium, has announced plans to build a data centre in neutral Switzerland, so that data collected in Europe will not be stored in an American facility, where it could be subpoenaed by the United States government.

In future the geography of the cloud is likely to get even more complex. "Virtualisation" technology already allows the software running on individual servers to be moved from one data centre to another, mainly for back-up reasons. One day soon, these "virtual machines" may migrate to wherever computing power is cheapest, or energy is greenest. Then computing will have become a true utility—and it will no longer be apt to talk of computing clouds, so much as of a computing atmosphere.