RADIO IN A CLOUD

By
Andrew P. Janitschek
Radio Free Asia
Washington DC

INTRODUCTION

Now that your station has mastered HD radio, you are ready to kick your feet up and watch the station run itself for a while, right? But wait; what is this? Cloud computing? Oh no, not another buzz word that management wants us to research! It will be no surprise if we have to produce another price analysis for the general manager. When will all these new ideas end? Well, never.

One of the many jobs of an engineer is to stay abreast of the latest technologies and where prudent, incorporate them in to every operational facet possible. Some say 'it is just not fair; we will never catch up' and they are most likely right. Ask any colleague and you will find they all share this sentiment too; it is impossible to stay abreast of all the latest technologies and trends, but we keep trying nonetheless. The world of broadcast engineering is one with continuing pressures to cut costs, maintain performances, and build flexible infrastructures. Beyond digitalization, convergence, and the greening of the broadcast industry, now on the horizon is cloud computing.

As engineers, we use the word 'cloud' as a metaphor for the Internet. We already represent the Internet as a cloud, as seen below, in broadcast designs, schematics and other drawings, but cloud computing does not relate to how we access the Internet so much as how we use the Internet. Cloud computing is coming of age and is ready to be our friend in broadcasting. While cloud computing is still not perfect, it can help you and I manage and maintain our broadcast station; for that matter, it can help in most any business environment. How? Business managers are increasingly looking to outsource their network infrastructure to managed service providers; there is no reason why we broadcasters cannot do likewise.

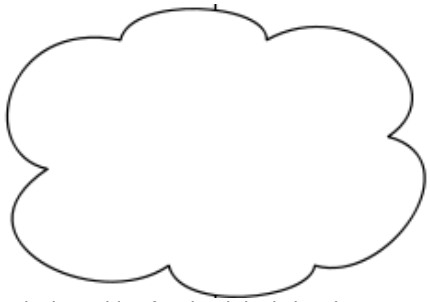

Typical graphic of a cloud depicting the Internet.

HEY! YOU! GET ON MY CLOUD!

The first step towards understanding how cloud computing can help us is to define it. One of the world's leading IT research and advisory companies, Gartner Inc, defines cloud computing as "a style of computing whose massively scalable and elastic, IT-related capabilities are provided 'as a service' to external customers using Internet technologies." Imagine if you will, terminals at every desk, but the common business applications the staff uses daily are delivered to them over the Internet while the software application is located on servers that could be most any place in the world; guaranteed they are most likely geographically separated from your broadcast station. For that matter, imagine all the updates to that software being handled off-site by a third party whom you contract with. Not only do they take care of managing these resources for you, but they also take care of maintaining and upgrading servers you use, all at a fixed cost; IT management is therefore overseen by you and others, but the 'grunt work' is not longer your worry because you now pay someone else to handle it. The two most common examples of cloud computing come from Amazon and Google; both rent their data-center resources to outside customers. We will examine each later.

Instead of buying or renting servers to deploy web applications, you rent computational power from a provider who already has servers with available capacity; to use the servers, you generally pay by the hour. Companies like Google and Amazon, with spare capacity on their own servers, will rent you time on their servers for you to use. By using a provider on the cloud, it is no longer necessary for you to buy expensive hardware or maintain it in your own data center; instead you use the computational units at the time you need them and then use as many as you need; that way you can handle temporary increases in demand for you applications. When the extra demand is gone, you stop paying for the use of the external servers as you stop using them. The added benefit is that you do not have to care about the hardware you are using, nor the replication of your data; through cloud computing, the system will maintain several copies of your data transparently so if your data is ever lost, you will be able to recover it.

Most cloud computing systems are built so that they are extremely reliable, that is, if any of the servers fail, the system will switch to a new server transparently. You,

the end user, will not even notice that the underlying hardware had a problem. The same happens for the data too. Advances like thus are what gives cloud computing its high reliability and though it's not perfect, researchers and developers are getting it there. In the near future, few web applications will experience downtime because of faulty hardware or problems in the datacenter. That means servers will be backed up and supported to an extent that even if a catastrophe destroyed a cloud provider's datacenter, systems will not go down. In this scenario, other datacenters around the world will take over and you, as a user, will still not notice a thing.

Companies have to be careful when deciding how many servers to rent from their cloud provider when moving away from traditional hosting services. Having too many servers could cost a lot more money in energy and computing power than needed and having too few could cause them to lose customers. Cloud hosting gives you more flexibility based on the resources you actually need at any moment in time. The thing to remember about cloud computing is that if some company can process or compute anything cheaper than you can, then it's more economic for you to outsource it to that cheaper provider; in its basic form, this gets down to the economics of running a business.

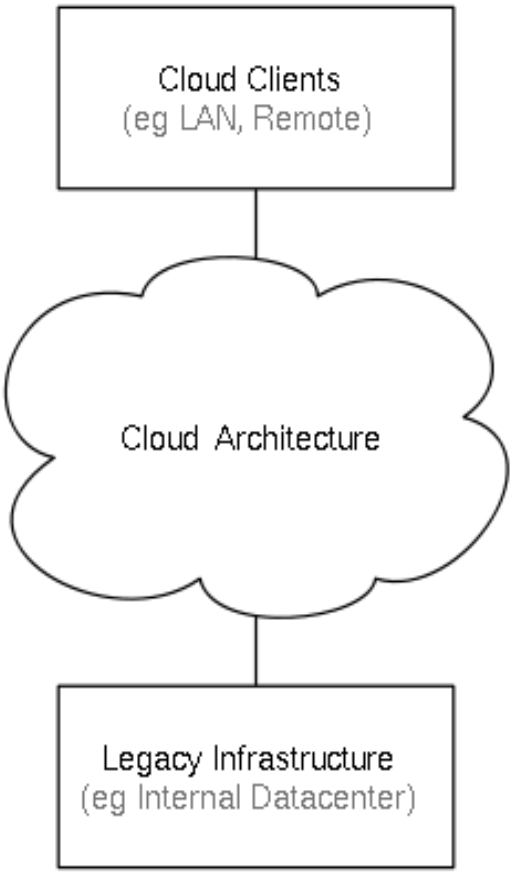

Typical cloud computing architecture.

SECURE THE CLOUD

What are some of the security issues? In the MIT's Technology Review, they published an interview in November 2009 with Whitfield Diffie who said, "A serious potential danger will be any laws intended to guarantee the ability of law enforcement to monitor computations that they suspect of supporting criminal activity. Back doors of this sort complicate security arrangements with two devastating consequences. Complexity is the enemy of security. Once Trojan horses are constructed, one can never be sure by whom they will be used." Whitfield Diffie is currently a visiting professor at Royal Holloway, University of London, and until recently worked at Sun Microsystems as a chief security officer. We will present more from Mr. Diffie a bit later.

Most everyone has an opinion on what is cloud computing. Another way to define cloud computing is to describe it as a means of accomplishing utility computing over the Internet using virtual servers. Cloud computing is, after all, a way to create new capacity without adding a large infrastructure. It can be the ability to rent a server or hundreds of them; it can be the ability to rent a virtual server, load software on it, turn it on and off at will, or clone it to meet short-notice increases in demand; or it can be the storage of huge amounts of data which is available to limited users and can be use applications which you authorize. Cloud computing can be supported by a provider that sets up a platform with the ability to scale automatically in response to changing workloads. Cloud computing can also the use of applications on the Internet that store and protect data while still providing a service; it can be using a storage cloud to hold various types of data; and it can provide usage of Web services.

By its concept, cloud computing is a change from the way business has handled proprietary data in the past. At a broadcast station, engineers and IT staff have traditionally installed, maintained and upgraded all software and hardware; all the while, data is stored internally, or possibly archived to an external location. With cloud computing, broadcast engineers no longer need to be experts with, or maintain total control over, the internal 'cloud' of their station. Through cloud computing, they can manage it all externally while maintaining seamless operations. In what has been described as a paradigm shift, cloud computing is a new supplement, consumption and delivery model for IT services based on the Internet, and typically involves the provision of dynamically scalable and virtual resources, as a service, using Internet connectivity.

Some of the important factors drawing businesses to cloud computing are its agility, cost savings, device and location independence, ability to serve multiple clients,

centralization, resource efficiencies, reliability, scalability, security, sustainability, and maintenance. When using cloud computing, agility improves with users being able to rapidly and inexpensively repurpose IT infrastructure resources. Overall costs can be dramatically reduced allowing capital expenditures to be diverted to operational costs. This helps lower barriers to cloud computing since the infrastructure is provided by a third-party and you do not need to purchase additional hardware for one-time or infrequent intensive computing requirements. Cloud computing by comparison is cheap. If you are scaling up an application, one cloud solution is Amazon EC2. Through Amazon, they will rent you computing power for as little as $.10 an hour; your time 'used' is only counted when the cloud is actually working on your behalf, so if you are not using it, it costs you nothing. In addition, at $.10 per gigabyte of bandwidth, this makes it extremely cheap to begin large scale growth.

Device and location independence in cloud computing lets users access systems using a web browser regardless of the geographical location of that system or what device they are using, whether it is a mobile phone, laptop, or some other equipment that works with the cloud. Since the infrastructure for cloud computing is typically off-site, accessed via the Internet, and provided by a third-party, users can connect from most anywhere. With an infrastructure that is large enough, third-party providers can also support multiple corporate clients and large numbers of users within one company, while serving each as well as all the others. This helps promote centralization of the infrastructure in areas where costs are lower for water, gas, electricity and real estate. Your IT staff does not do not need to design your internal infrastructure to match the capacity for the highest possible load either; with cloud computing, the only consideration is that your peak-load should easily be handled by your third-party provider. Lastly, when the cloud computing infrastructure is large enough, there are also improvements in the efficiency of systems that often used only 10-20% of their capacity in the past.

Reliability improves with cloud computing through the use of multiple redundant sites; this makes working with the cloud suitable for business continuity and disaster recovery. Nevertheless, as cloud computing has grown, many major cloud computing providers have suffered outages leaving IT and business managers with little to no ability to fix a problem when they are affected.

The scalability of cloud computing is excellent when using dynamic sharing of resources on a self-service basis without users having to plan or schedule for peak loads in this sense the cloud is infinitely scalable. It is not necessary to worry about clustering, nodes, or serving up content from one part of the globe to another. Cloud computing lets you build as much capacity and bandwidth as you're willing to pay for.

When incorporating cloud computing into business models, security improves due to centralization of data, increased resources dedicated towards security, and because your cloud provider already has a capable cloud support staff standing by to help if needed. Security is often better than traditional systems mostly because third-party providers are able to devote more resources to security that many of their clients cannot afford. Providers typically log accesses, but accessing the activity logs themselves requires specific training and security is further increased when data is spread out over a wide area and across a number of devices.

Sustainability comes about through improved resource utilization and more efficient systems. Add to this that fewer servers and computers are needed and this contributes to a reduction in energy usage. While sustainability is improved, so it maintenance; cloud computing applications are easier to maintain since they do not have to be installed on every user's computer. This makes them much easier to support, update and upgrade since changes reach each client immediately.

The best way to understand the force behind cloud computing and make it accessible to everyone is to look at it in financial terms. For the sake of argument, let us assume you are an engineer at a large broadcasting station and you want to launch a new project. Under normal circumstances, you will likely have had to write out a purchase requisition to buy the new hardware and software necessary. Your purchase requisition must now go through the usual approval cycle, and once it has been signed, you will most likely need a 2-3 month order timeframe. Now you have already been delayed 3-4 months before your project can really gather steam and move quickly along. With cloud computing though, you could log on to some provider, like Amazon's Elastic Compute Cloud (EC2), and purchase capacity for as little as a few cents an hour. Instead of spending up to $50,000 or more for hardware and software, you can be up and running instantly and with a completed solution for just a few hundred dollars. This does not even address the normal delays of weeks or months if you were to try to build this infrastructure from scratch or expand upon your existing capacity. This is the main sales pitch from many third party cloud computing service providers. Obviously, they make sense. From the point of view of any broadcast engineer, this is really a great idea. It gives us the flexibility we need and takes away the headache of building and maintaining the infrastructure. The entire infrastructure resides with a third party who already has huge economies of scale and can then pass those saving on to you.

Your applications working with Amazon EC2

Amazon has very publicly and aggressively run with this idea. No doubt, we will also see other major cloud computing providers like Cisco, IBM and Google market their services assertively to all major industries. With their deep pockets and seemingly endless access to hardware, software and numerous data centers, they will not have problems providing any level of service to any organization, no matter how large or how small, and can offer their services at great prices. This is the concept known as Infrastructure as a Service (IaaS).

Besides IaaS, cloud computing also embraces Saas and Platform as a Service (PaaS). In these two scenarios, companies offering either a comprehensive solution stack as a service, PaaS, or online use of discrete software applications on a per-use basis, SaaS. A model like this makes a great deal of sense. There are also problems with this too. The largest problem that should concern you and others with cloud computing is security. No one is questioning or doubting the ability of any third-party providers to continually work on stabilizing their infrastructure and platforms but instead, question whether or not potential customers genuinely trust them enough to put their critical data on the cloud where they may or may not be able to access it when they most need it. It may be irrational to fear this loss of control, but it is only natural that managers and engineers should feel this way. To help allay these fears, many companies and cloud providers are researching and developing security overlays that sit on top of the public cloud. There's also a movement toward private clouds, where a large company might build its own cloud so as many as thousands of employees can plug into it. This would help alleviate fears of losing control of data or processes.

Cloud computing poses several security risks. Not only could stored data be stolen by hackers or lost to due to breakdowns, but a third-party provider might mishandle data, or be forced to turn it over due to pressure from a government or to a subpoena. Security breaches like this have already happened too. In 2008, a single corrupted bit in messages between servers used by Amazon's Simple Storage Service (S3), which provides online data storage, shut the system down for hours; in 2009, a hacker who correctly guessed the answer to a Twitter employee's personal e-mail security question was able to grab all the documents in the employee's Google Apps account; then there is the bug that compromised the sharing restrictions placed on some documents in Google Docs so that differences were wiped out and anyone with whom you shared documents could see all the documents shared with anyone else. Another example of security problems with cloud computing took place just this past October when over a million T-Mobile Sidekick smart phones lost data after a server failed at Microsoft's subsidiary, ironically called Danger. Luckily, most of the data was recoverable. Even with these recent examples, the proponents of cloud computing insist the clouds are more secure than anything else we are using now. In a way, their argument makes sense; cloud providers stay on top of security threats more effectively than millions of individuals and thousands of companies running their own computers and server rooms. Google itself points out that for all the publicity surrounding the problems with Google Docs, less than .05 percent of all documents hosted by Google have been affected. This is good for those who did not experience any problems; not very comforting though for those who did run into problems. Given the industry's rapid growth, the murkiness of its current security standards, and the anecdotal accounts of breakdowns, it's not surprising that many companies still balk at the idea of putting their proprietary data on the cloud. Though security is pretty good, cloud providers still need to prove their reliability over the long term.

Cloud providers do not have a virtual steel fence to sell you yet, but they can promise to keep your data on servers in specific countries, like the United States, or economic areas for regulatory compliance or other reasons. They are also working on virtual walls. In August, Amazon announced plans to offer a "private cloud" service that ensures more secure passage of data from a corporate network to Amazon's servers. All the while, new technologies in security are still being developed. For example, Microsoft suggested a way to prevent users of one virtual machine from gleaning information by monitoring the use of shared cache memory by another virtual machine on the same server. In another example, IBM proposed a new security mechanism that will examine new virtual machines as they enter the cloud. Software would monitor each one to see how it operates and ensure its integrity, in part by exploring its code. These technologies could be ready as soon as 2011. Emerging encryption technologies, however, could protect data in clouds even as users

search it, retrieve it, and perform calculations on it. This could make cloud computing far more attractive to industries such as banking and health care, which need security for sensitive client and patient data. In a broadcast network, servers contain client information, costs for spots, rotation of playbacks, readers, payroll, music, and more need security if stored on the cloud. For starters, several research groups have developed ways of using hierarchical encryption to provide different levels of access to encrypted cloud data. A patient, for example, could hold a master key to his or her own electronic medical records; physicians, insurers, and others could be granted sub-keys providing access to certain parts of that information.

Ideally, we should make it easier to work with sensitive data that needs to be encrypted so that unintended viewers cannot see it if it were exposed by a hack or a glitch at the cloud provider. To find and retrieve encrypted documents, groups at Carnegie Mellon University, the University of California, Berkeley, and elsewhere are working on new search strategies that start by tagging encrypted cloud-based files with encrypted metadata. In order to perform a search, the user encrypts search strings using mathematical functions that enable strings to find matches in the encrypted metadata. No one, not even a hacker, could see the document or even the search term used.

If cloud computing does become secure enough to be used to its full potential, new and troubling issues may arise. For one thing, even clouds that are safe from ordinary hackers could become central points of Internet control. Regulators, courts, or overreaching government officials might see them as convenient places to regulate and censor. Cloud providers themselves could crack down on clients if, say, copyright holders apply pressure to stop the use of file-sharing software. What worries some people the most is the ability of a government to get access to private data, fewer constitutional protections, increased government censorship, and increased vendor, or government, controls on innovation. But such concerns aren't stopping the ascendance of the cloud. And if cloud security is achieved, the benefits could be staggering.

WHO IS IN THE CLOUD

Cloud computing is already changing the way we use computers; Gmail, Twitter, and Facebook are all cloud applications. Web-based services like Amazon, Google, and Rackspace are drawing a number of customers because of their low pricing and overall efficiency. Rackspace delivers enterprise-level hosting services to businesses of all sizes and kinds around the world. Google's Chrome is meant to provide easy access to cloud applications. Some of Amazon's clients for cloud services include the New York Times and pharmaceutical giant, Pfizer. Even government agencies are joining the cloud. For example, the City of Los Angeles uses Google's Apps for e-mail and other routine applications, and the White House recently launched a website to encourage other federal agencies to use the cloud; that website is www.apps.gov. Airline, marketing, retail, and financial institutions are also benefiting from the use of cloud computing.

Nothing in life is every perfect and certainly the same applies to cloud computing. For those that are frugally minded, lower cost cloud computing solutions, like Amazon S3, you do tend to have greater chances for latency, so let the buyer beware and sometimes you also get what you pay for. For lower cost solutions, the network connection may not be fast enough to serve rich media at scale. Many services that are available, including WordPress.com, use Amazon S3 as a "cold cache." Cold cache means that this is the last location content is served from, and then, only if needed. Depending on the application, the latency of the network may make it impossible to cold cache every page load.

WordPress.Com

Lastly, cloud reliance can cause significant problems if you must surrender control of downtime and outages. Over the past two years, Amazon has had significant outages; in one case it lasted 8 hours. Reliance on the cloud can be a real headache when time is money. On Wednesday, December 9, 2009 Amazon Web Services had a power failure at its main data center in northern Virginia. Amazon reported an "underlying power issue" that affected its US-EAST-1 availability zone (Amazon has four availability zones in the U.S.; two on the East Coast and two on the West Coast). Amazon did not immediately confirm a weather storm the previous day was the cause but the same datacenter suffered a similar outage in July 2009 when it was struck by lightning.

As defined by InfoWorld, software as a service (SaaS) in cloud computing delivers a single application through a browser to thousands of customers using a multi-client architecture. On the customer side, it means no upfront investment in servers or software licensing; on the provider side, it means that it is just one more application to maintain. Costs for the provider are low compared to conventional hosting; one of the best

known enterprise applications is Salesforce.com. Salesforce is one of the pioneers of using the SaaS model for distributing business software. In this case, access to the software is purchased on a subscription basis and all the while hosted offsite. Salesforce is best known for their Customer Relationship Management (CRM) products, which it delivers to businesses over the Internet using the SaaS model. SaaS is also common for Human Resource applications and has even worked its way up the food chain to Enterprise Resource Planning (ERP). Just a few years ago, no one would have been able to foresee the quick acceptance of desktop applications like Google Apps and Zoho Office.

In an enterprise sense, cloud computing is often handled by Content Delivery Networks (CDNs). Some of the better know CDNs are Akamai, Panther and Limelight; Radio Free Asia uses Akamai. These enterprise solutions can be as cheap as $0.25 per 1GB of bandwidth and up to as much as $1 per 1GB; less expensive solutions include Amazon's EC2 and Google App Engine.

Broadcast companies looking to expand their offerings with web apps, downloadable materials or provide a service relying heavily on rich media such as images, HD video, or streaming audio should consider the Cloud as a viable scaling solution.

Other key issues in cloud computing include the need for standards to measure and ensure the quality of managed services, the role of governments in regulating the industry, as well as what to do during emergencies. One of the main benefits of managed services is the savings on capital expenditures. A capital expenditure is when business spends money either to buy fixed assets or to add to the value of existing fixed assets. Most organizations should weigh the total cost then decide if it makes sense to in-source, outsource, or gradually move to managed services. Costs to consider are the asset lifecycles, return on investment, core business technologies, security, maintenance, manpower and space. Many enterprise customers are shifting toward faster 'time to market' processes in order to keep their competitive edge. While there is a great deal of optimism about cloud computing, we must realize there are significant hurdles that managed service providers need to overcome before it becomes the primary industrial platform. The largest single obstacle is the lack of cloud computing standards.

CLOUD – BUY IT NOW

In an IT industry first, Amazon has launched a spot market to sell off excess capacity on its EC2 platform. Amazon has begun a beta version of Spot Instances, which lets customers bid for and buy capacity at prices set by demand and availability. Amazon's CTO Werner Vogels said "as far as I know, [this] has no equivalent in the IT industry." He said this market was made possible by the enormous economies of scale with Amazon's infrastructure services. Additionally, he said that customer-managed pricing would "make new areas of computing feasible for which the economics were previously unfavorable … customers bid any price they like on unused Amazon EC2 capacity and run those instances for as long their bid exceeds the current Spot Price." The EBay-like auction pricing gives clients tight control over the maximum cost for their workloads. The auction platform is best suited to customers with non-mission critical workloads on a flexible timetable; these could be small, nebulous processes or even huge undertakings like media conversions or web crawling for broadcasters or other companies.

International Data Corporation predicts business IT spending on cloud services will increase from $16 billion in 2008 to $42 billion by 2012. When Amazon introduced the Elastic Compute Cloud (EC2) in 2006 it was a turning-point in the effort to change computing into a utility, like electricity. Suddenly, everyone could scroll through an online menu and hire as much computer power as they needed while paying for it at a set rate. It initially cost 10 International Data Corporation predicts business IT spending on cloud services will increase from $16 billion in 2008 to $42 billion by 2012.

When Amazon introduced the Elastic Compute Cloud (EC2) in 2006 it was a turning-point in the effort to change computing into a utility, like electricity. Suddenly, everyone could scroll through an online menu and hire as much computer power as they needed while paying for it at a set rate. It initially cost 10 cents

per hour to use Linux. Two years later, Microsoft Windows became available at the price of 12.5 cents per hour. Both systems run on virtual machines that can be created and configured in an instant then also disappear when no longer needed. As needs grow, Amazon's customers could now increase their computing capacity simply by purchasing more computing power available on the cloud while Amazon would takes care of maintaining the data center and network. Naturally, the virtual machines are running inside the thousands of servers clustered in Amazon's global data centers. The service is efficient, cheap, and accessible to everyone; private individuals, companies, universities, research laboratories, and government agencies. But it also poses a potential threat. EC2 brought us all something once confined mainly to corporate IT systems: engineering in programs called hypervisors create and control virtual processors, networks, and disk drives, many of which may operate on the same physical servers. Computer security researchers had previously shown that when two programs are running simultaneously on the same operating system, an attacker can steal data by using an eavesdropping program to analyze the way those programs share memory space. They surmised that the same kinds of attacks might also work in clouds when different virtual machines run on the same server. It seemed though that the possibility of a hacker finding their intended victim on a specific server seemed remote. Nonetheless, in 2009, three computer scientists at the University of California, San Diego, and one at MIT went ahead and did it. They hired some virtual machines to serve as targets and others to serve as attackers and then they tried to get both groups hosted on the same servers at Amazon's data centers. They succeeded in placing malicious virtual machines on the same servers as targets 40 percent of the time, all for a few dollars. While no one actually stole any data, the researchers said that this was proof that stealing data using this system is theoretically possible. They also demonstrated how the advantages of cloud computing could create new ways of bypassing security measures, though Amazon stresses that no one has successfully attacked EC2 in this manner and that Amazon can prevent attacks like this now and in the future.

In his online blog on ZDNet.com, Phil Wainewright wrote, "If I read one more article about what MIT Technology Review in its January [2009] lead story is calling "*the security problem inherent in the size and structure of clouds*," then I swear I am going to burst a blood vessel." Mr. Wainewrite goes on to write tongue-in-cheek, "[how the people at MIT] seem to believe we should all be very scared" because stealing of data stored from the cloud is theoretically possible; "Oh my goodness, how awful!" His view is one of many that feel that worrying about the loss of data should receive no more attention than worrying about someone breaking into a facility and stealing data from filing cabinets; the possibility is always there but vigilance is the best method of preventing data loss, not ignorance of the technology. Mr. Wainewrite ends his blog by quoting Mr. Diffie, "The advent of radio posed similar issues a century ago … Radio was so much more flexible and powerful than what it replaced — the telegraph — that you had to adopt it to survive in business or war. The catch was that radio can be picked up by anyone. In radio's case, fast, automated encryption and decryption technologies replaced slow human encoders, making it secure enough to realize its promise. Clouds will experience a similar evolution."

IS IT A CLOUD OR A CDN?

A content delivery network or content distribution network (CDN) is a system of computers containing copies of data, placed at various points in a network in order to maximize bandwidth usage for accessing to the data. A client accesses a copy of the data near to their location as opposed to everyone accessing the same central server; this helps avoid bottlenecks at that server. Content types include web objects, downloadable objects like media files and documents, applications, real time media streams, and other components of internet delivery. Amazon Virtual Private Cloud (Amazon VPC) is a secure and seamless bridge between a company's existing IT infrastructure and the Amazon Web Services (AWS) cloud. Amazon Virtual Private Cloud (VPC) enables enterprises to connect their existing infrastructure to a set of isolated Amazon Web Services (AWS) resources using a Virtual Private Network (VPN) connection. The goal is to expand their existing security applications like firewalls and intrusion detection systems to include their AWS resources. Amazon VPC integrates today with Amazon Elastic Compute Cloud Amazon EC2 and integrates with other AWS services. Amazon EC2 is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. As with all Amazon Web Services, there are no long-term contracts, minimum spending or up-front investments needed; the same with Amazon VPC, you pay only for the resources you use.

CLOUD BY AMAZON

Amazon CloudFront was released at the end of 2008 and is the AWS for content delivery. It integrates seamlessly with Amazon S3 to provide fast distribution of content with high data transfer speeds using the cloud. Like other Amazon Web Services, it requires no upfront commitments and is a pay-as-you-go service. Amazon CloudFront is designed to be fast; it will cache

copies of the content in locations close to the end-user's location, significantly lowering content access latency. CloudFront will maintain highly sustainable data transfer rates even when you are distributing larger objects. Amazon CloudFront is useful for many different application scenarios such as giving customers low-latency access to popular objects and by protecting your site from surges in usage. Other examples are low-cost delivery of rich media and sustainable fast transfer rates for software distribution.

As of now, Amazon customers achieve scalability, reliability and cost-effectiveness through the use of Amazon S3 and can easily integrate with Amazon EC2. Amazon CloudFront builds further on that integration by making it easier to distribute content world-wide. Amazon S3 provides durable storage of your data and their network of edge locations in Europe, Asia and the United States will deliver the content to your customers with low latency from the most appropriate location. Here are the edge locations:

- United States: Ashburn (Virginia), Dallas/Fort Worth, Los Angeles, Miami, Newark, Palo Alto, Seattle and St. Louis
- Europe: Amsterdam, Dublin, Frankfurt and London
- Asia: Hong Kong and Tokyo

These edge locations work together to direct customers' requests to the edge location that can provide the response with the lowest latency.

GOOGLE CLOUD FOR FREE

Google App Engine is a cloud based platform that enables Google's customers to run web applications on Google's architecture. Basic usage of the Google App Engine is free. After signing up for a free account on Google Accounts, applications deployed to the Google App Engine may use up to 500 MB of storage, 10GB incoming bandwidth per day, 10 GB outgoing bandwidth per day, 200 million megacycles of CPU usage per day which is enough for about 5-million page views per month and 2000 emails per day. If these levels are exceeded, Google offers consumption-based pricing. As long as you do not go over the limitations of the Google App Engine, you costs for usage are nil. In the current version of the App Engine toolkit, applications must be implemented in Python. The runtime environment includes the full Python language and most of the Python standard library.

Google apps that work with the cloud.

CLOUD FOR SALE

One of the fears of using a free application, like Google's App Engine, is that innate fear that 'we get what we pay for.' There are many cloud computing services available that offer different prices and different structures. Here are a few worth your time to research:

- 3Tera – AppLogic. Grid OS used as cloud computing platform by service providers and enterprises.
- Agathon Group. They are a cloud provider. Services include highly available VPS, virtual private datacenters and ready-to-use LAMP stacks along with self service ordering. Custom development and managed services available.
- Appistry. Cloud computing middleware. Enables easily scalable cloud computing.
- Bungee Connect. Provides end to end tools and systems required to develop, deploy and host web applications.
- CloudStatus and CloudClimate. They are both cloud enablers providing real-time

performance trending of cloud infrastructures. See below for more on CloudClimate.

- Coherence. Know for the Oracle Coherence Data Grid for EC2 and other cloud platforms.
- Concentric. The Concentric Managed Servers offer high performance, secure, reliable hosting without the upfront IT infrastructure costs.
- Enomaly Inc. Service provider and cloud enabler. Developer of the Enomalism Elastic Computing Platform & Elastic Drive.
- Flexiscale. Another supplier of instant web servers with some advanced features like auto-scaling coming soon.
- GigaSpaces. Middleware for the cloud, or "cloudware."
- GoGrid. Instant, on-demand servers offering "control in the cloud". Deploy Windows/Linux servers via web-interface in minutes.
- Layered Technologies. Cloud provider of on-demand hosting and cloud and utility computing solutions through its brand GridLayer.
- Kaavo's IMOD. An easy to use online application; cloud computing made easy.
- Mosso. Rackspace's cloud hosting service.
- Nirvanix. Cloud storage.
- RightScale. Another cloud enabler.

WATCHING THE CLOUDS

As mentioned in the list of providers above, CloudClimate provides real-time performance trending of cloud infrastructures. They display monitoring results from a globally distributed installation of PRTG Network Monitor, a network monitoring software from Paessler AG. PRTG works with one core server installation and a number of remote probes used to measure system performance and to remotely monitor performance of network services. The remote probes for CloudClimate are installed on virtual systems in the hosting clouds below with more coming soon:

- Amazon EC2 US East Region
- Amazon EC2 Europe West Region (Ireland)
- GoGrid Cloud Servers (San Francisco)
- NewServers.com (Miami)

They are also running a number of probes on selected low-cost servers hosted by:

- VPSLand (Atlanta)
- HostEurope (Cologne, Germany)
- Webhosting.co.uk (London)
- Usonyx (Singapore)
- HostingPanama

CloudClimate monitors the internal performance of cloud servers by running short load tests for CPU, memory and disks every 5 minutes. All probe systems send HTTP and PING requests to each other every minute to measure network performance. The results are shown in the graphs on their website, **www.CloudClimate.com**. The following graphic shows the detailed monitoring results for a "m1.small" instance running in the US West zone of Amazon EC2. CloudClimate measures the performance of the system itself using short load tests for CPU, memory and disks.

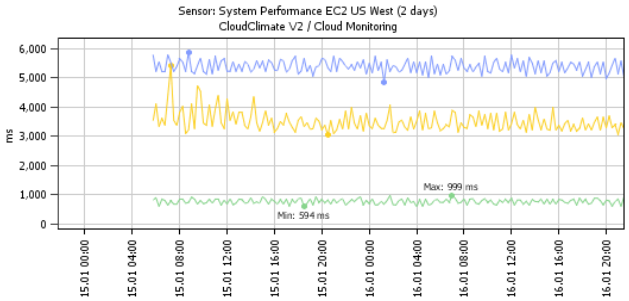

Jan 16, 10 monitoring results for the US West zone of Amazon EC2

GLIDING THROUGH A CLOUD

At this year's Consumer Electronics Show (CES), Glide launched a powerful operating system for the web browser, Google Chrome browser. Glide changes Google Chrome from a standard browser into an operating system with a rich application suite and 20GBs of free storage. In her January 2010 Cloud Computing Journal article, Ms. Yeshim Deniz writes, Glide pumps up "Google Chrome with file synchronization and automated file format translation to speak with multiple remote Windows, Mac and Linux desktops and all mobile platforms, web content capture, a rights based file management system to manage data across desktops and mobile devices, a word processor, presentation app, spreadsheet, photo editor, email client, online meeting app, drawing tool, web publishing app, contact manager, calendar and much more." As of this writing, Glide is available for free at **www.GlideLife.com**.

SAFETY CLOUDS

Cloud computing is not just for standard business; it has already moved into the area of public safety and is doing quite well as companies introduce software as a service (SaaS) for communications between fire, police and emergency agencies. Cloud computing has made progress with state and local public safety communications with the rollout of a SaaS solution that links radio systems. In an article published in April 2009 by Federal Computer Week, Alice Lipowicz details how Twisted Pair Solutions of Seattle has partnered with Fpweb.net, a managed hosting services

provider, to provide IaaS solutions to state and local agencies powered by cloud computing. Tasks are moved off agency servers and onto shared data centers. For public safety, the goal is to combine cloud computing with voice-over-IP radio systems to give differing radio systems and networks the ability to work with each other. Twisted Pair provides a cost-effective solution when compared with hardware-enabled interoperability. Through these efforts, public safety agencies currently can buy bridging devices that connect different radio systems, build advanced radio systems, and use cell phones or satellite phones as backups. This is a boon to public safety as fire, police, emergency management and other public agencies can now connect their private push-to-talk radio systems to each other. This not only gives everyone in a single agency access to their communications system, but also to in other agencies that must work with that system. In one example, the Department of Homeland Security (DHS) recently approved the Bridging Systems Interface, a VOIP specification for bridging radio systems. Twisted Pair said its solution also uses the bridging systems interface.

HERE TO SAVE BROADCASTING?

Broadcasting and other communications companies are asking how they can shift their business models so that they may emerge as one of the winners in the new cloud computing world. Communications companies are looking at everything from opening their own data centers to hosting other third parties, to offering their own applications and offering them using PaaS. For example, a service provider might offer its billing system as a service in the cloud. Customers could pay a 'fee per-transaction' or on another revenue-sharing basis and plug into the billing system. Other opportunities surround their service delivery systems, their location systems and so on. This goes back to the idea of a "two-sided" business model where traditional communications companies change from not only a company that delivers services but also to opening their core capabilities and offering them to customers. In this case, the services would be offered through the cloud. It is certainly nothing trivial to take your internal systems and reconfigure them so they are suitable for consumption by paying public. Risk and scalability are the two things to watch for here. On one hand, we could make a significant investment to open typically closed internal systems, only to find that there is no market for such a platform. On the other hand, we could find the service a resounding success and realize that we simply cannot handle the global scale that cloud computing may demand of us.

"We are developing bidirectional solutions to communicate more effectively with the audience."
- David Baden, CTO, Radio Free Asia

CLOUDY MEDIA

While the concept is still new to most broadcasters, some are already stepping into the fray and adopting cloud computing measures to reduce costs and to provide better services to their audience. Without a doubt, radio and television stations handle large amounts of data. Many of us may not think of our broadcast station as a data center, but in truth, it is just that; if you are not going right to the cloud then, there are data center considerations that can be applied to building or expanding your station's IT facilities.

In November 2008, Mark Welte's article, Design Best Practices for Data Centers, was published by Radio magazine. In that article, Mr. Welte writes, "Data centers are going through revolutionary changes, due to changing market conditions and technological demands; these include the inability to secure large sums of capital dollars to fund data center projects, the high cost of energy, and the pursuit of a "low carbon diet," in the wake of recent cap-and-trade legislation." Other issues changing the broadcasting IT landscape are innovative information technologies like cloud computing and virtualization. They are forcing engineers to rethink the design, operation and planning of a data center. Mr. Welte writes, "On the other side, the data center containerization industry (DCCI) aims to have the capital investment associated with power and cooling infrastructure be a direct match with computing. Even the UPS industry is in competition for the most efficient and scalable products, with containerized power and cooling solutions to consider. Having a firm design foundation that can adapt to change is paramount to a facility's success."

It is important to carefully plan the work and to continually work with the plan. Consult with design teams experienced in the design of your type of broadcast facilities. If you do not know anyone, ask around; your friends and colleagues in the industry are usually the best source for companies and consultants that will help you accomplish your goals. If you blow off the need for getting advice from experts, you certainly will have bragging rights to everything that goes right, but you will also have to swallow a bitter pill for everything that goes wrong. Initial designs are too often not fully developed or are too vague. Designs that are thought through and specific will start with a

narrative of the project's background and provide details of required objectives, and specify the goals. Performance and maintenance criteria must also be addressed while clearly defining the expected lifespan and overall quality of the project.

In order to guarantee continuous operations with as few interruptions as possible, there must be an emphasis on the HVAC system and energy usage; and in this age, keeping it green is even more important. In my years in broadcasting I have worked in a few stations where the HVAC system never worked correctly. Thankfully, there are many engineering tools that can help you design and forecast performance behaviors. For example, today's CAD programs not only help you with high-end designs, but also assist in solving structural issues, power usage, safety and more. Designing in 3-D is no longer just a 'cool aspect' of CAD designing, but is becoming more of a must! You can now render your data center in 3D and literally take upper management on a virtual walk-through without ever starting any construction. As an aid to CAD designers, Radio Free Asia provides hundreds of free 3D drawings through our Broadcast Open Development Exchange Initiative at **www.TechWeb.rfa.org**; look for 3D Project link. The 3-D Project serves as a download site for the free exchange clearinghouse for 3-D tutorials, 3-D object drawings and material/texture bitmaps. This site also includes detailed project standards information for file type, and the file/layer naming conventions.

Now, power analytics, a robust set of software tools that designs electrical systems can be used to perform predictive analysis and what-if scenarios. One set of power analytics tools, Paladin Design Base and Paladin Live from EDSA, are used in demanding customer sites to simulate and analyze systems rom a variety of static or dynamic perspectives. A screenshot from the EDSA software is below. They include the ability to model and embed detailed control logic of intelligent electronic devices that control power flow throughout the system. Power analytics systems act as an on-board electrical power system expert, to intelligently filter the power system sensory data, help owner/operators understand the real-time health of their electrical power system, as well as diagnose whether the system is stable, deteriorating, or becoming overloaded. This helps us make decisions about electrical power long before problems begin.

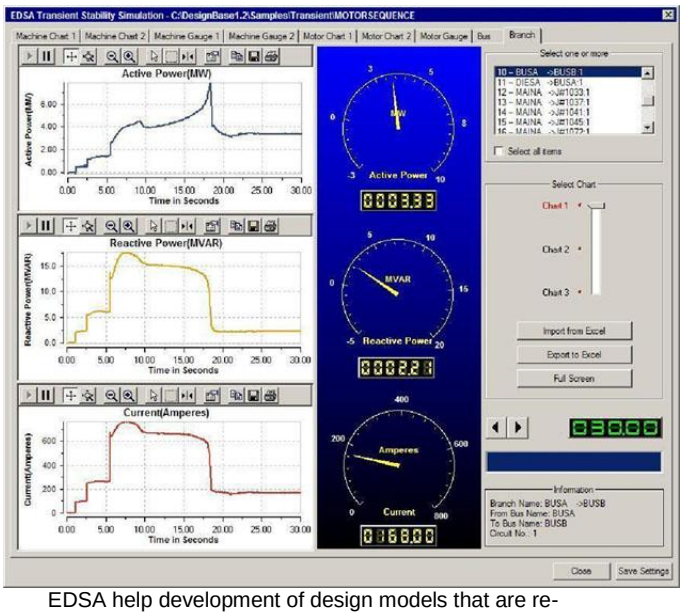

EDSA help development of design models that are re-deployable in real-time.

For projects with a life expectancy over 10 years, the operating costs of the mechanical system will likely exceed the initial installation cost. The choice of which HVAC system to purchase initially will also determine options for future upgrades and expansions. If you are using blade servers, you need to know they consume 10 times more total power and need at least four times as much cooling capacity. According to an Uptime Institute study, power consumption in data centers has increased over 600 percent in just seven years. As the graph below from ASHRAE demonstrates the trend for increased heat load will increase exponentially in the future.

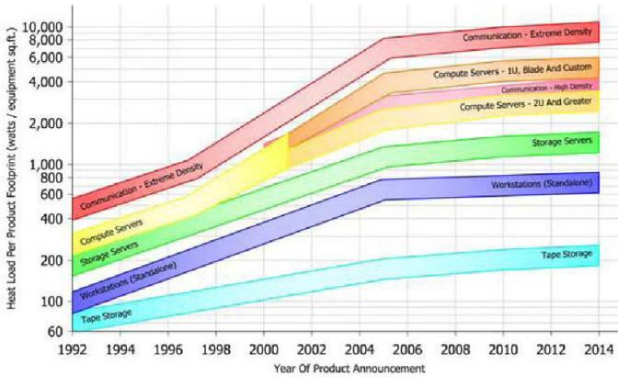

ASHRAE shows increasing datacenter power trends and cooling applications.

There is more to designing the best possible facility, but understanding the energy and cooling systems will help ensure a successful design. As Mark Welte wrote, "Armed with a solid design plan and incorporating new software tools that can help design and perform critical what-if scenarios, achieving more energy efficiency and higher facility uptime can now be quickly achieved. With these thoughts in mind, maybe going to a cloud

computing system instead of building a data center is not only faster, but cheaper. ”

THE SUN IN CHYRON’s CLOUD

The move towards new workflows in broadcasting is gaining pace. With cloud computing, it could mean doing away with expensive hardware on site. If you thought cloud computing is still years away, it might be time to rethink your position on that.

AXIS is a news graphics application from Chyron giving you access to high definition graphics from anywhere in the world without the need to invest in your own expensive system. According to Chyron, “AXIS is a CSP (Cloud Service Provider), meaning that you do not have to worry about installing any expensive proprietary equipment to make it work. All you need to do is subscribe to the service and we handle all of the equipment and infrastructure headaches that plague traditional solutions.” Chyron believes AXIS, which is a SaaS, helps define new workflows. It is often what is called a disruptive technology in that it is changing the rules in our broadcast industry. It is creating new ways of doing things. AXIS not only offers cost savings, but also unique solutions and unexpected benefits.

Chyron is pushing to sell cloud services.

In August 2009 Broadcast Engineering announced that KLAS-TV in Las Vegas needed a way to sustain newscast quality while reassigning the resources needed for creating news graphics, Chyron’s AXIS is a SaaS package that works “over the Internet providing high-quality graphics for broadcast and Web-based video outlets. Requiring no hardware or software, AXIS enables anyone with an Internet connection to fulfill instantly the need for broadcast-quality news graphics, still or animated maps, bar and pie charts, financial stock quotes, and weather graphics. According to KLAS creative director Lee Minard, the primary use of AXIS is for news, but it also supports some of the station’s weather requirements.”

KLAS-TV news graphic using Chyron’s AXIS online content creation system.

SaaS is the best example of a disruptive technology in broadcasting since the nonlinear editor. The television production company, Endemol, is already starting to work with the cloud by using the Internet to bring together reality shows from around the world for the social network site, Bebo.com. The philosophy being cloud computing works well for Endemol’s type of television production; instead of having a traditional film crew traveling across the world, the new model of television production can be done by one person who edits and uploads their material to the Internet rather than working on a standalone system in some backroom someplace.

Endemol is pushing for more cloud production.

Orca is a proof of concept application Siemens developed to test the cloud in a broadcast environment. It was a success, but the development team also said the broadcast industry must be realistic. The bandwidth needed for media production must be huge! Since movie files can be as much as 5TB, there is no easy to move them in and out of the cloud yet, but smaller scaled productions and different types of applications are very relevant and very applicable to using cloud computing.

FAIR TO PARTLY CLOUDY FOR NOW

Cloud computing already offers web-based enterprises the opportunity to deliver business solutions without the need for local hardware and software platforms. Can online services transform the way media is created and delivered? Can the same model be applied to broadcasters needing a great deal of bandwidth and processor-intensive post production? Can broadcasting facilities be freed from the constraints of their geographical location? Can the cloud provide quality of service we regard as essential for production and delivery to today's critical audience? Cloud computing is still in its infancy within the broadcasting industry but its benefits, especially the financial benefits, mean that cloud computing is set to grow quickly.

When bandwidth and connectivity is no longer a constraint, the rules change.
- IBC2009