



WHITE PAPER

An Introduction to
Cloud Computing
in the Federal
Public Sector

AUGUST 2010



The intended audience of this document comprises senior technical executives, architects, engineers, and developers who need to understand the concepts of Cloud Computing and whether they are applicable to their unique requirements.

An Introduction to Cloud Computing in the Federal Public Sector

The Cloud Computing industry represents a large ecosystem of many models, vendors, and market niches¹ each with their own approaches to exploiting the gains that cloud promises. In May 2010, a simple Web search of the term "Cloud Computing" returned over 14 million hits, creating an intractable task for any leader, team, or agency trying to define the cloud and adapt it to further their mission. This vendor-neutral white paper presents the concepts and application of Cloud Computing in the U.S. public sector and can be used as a starting point for formulating an agency's cloud strategy.

The economic gains of Cloud Computing are so compelling that OMB budget and guidance contains cloud requirements as follows²:

- In FY10, OMB requires agencies to launch a series of Cloud Computing pilots across the government using the E-Government Fund.
- In FY11, the OMB will require agencies to develop an alternative analysis discussing how they could use Cloud Computing for all major technology projects in FY12.
- In FY12, OMB will require agencies to provide complete alternatives analysis for IT programs that are new development or are steady state, explaining how they could incorporate Cloud Computing in their environment.

Contents

The Five Essential Characteristics of Cloud Computing	1
Cloud Computing Service Models	2
Cloud Computing Deployment Models	4
Cloud Computing Use Case Scenarios	5
Security in Cloud Computing.....	6
FedRAMP	7
Conclusion	8

The Five Essential Characteristics of Cloud Computing

There are over 20 working and broadly accepted definitions of Cloud Computing. In the public sector, the definition of Cloud Computing is established by the National Institute of Standards and Technology (NIST) Computer Security Division. Although brief in length, the definition presents five essential characteristics, three service models, and four deployment models, all of which are important to understanding the nature and terminology of the cloud, which is:

"... a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."³

Cloud Computing is a form of delivering IT services that takes advantage of several broad evolutionary trends in IT, including the use of virtualization; the decreased cost and increased speed of networked communications, such as the Internet; and overall increases in computing power. As such, any definition of Cloud Computing will be somewhat broad and subject to interpretation. [U.S. Government Accountability Office Report to Congressional Requesters, Information Security: Federal Guidance needed to Address Control Issues with Implementing Cloud Computing, May 2010, <http://www.gao.gov/new.items/d10513.pdf>]

Within the ecosystem, vendors frequently co-opt the five essential characteristics to apply the "cloud" label to their

service or offering, leading to widespread confusion and misunderstanding about whether a particular software, platform, or infrastructure should be called "Cloud Computing." By labeling the characteristics "essential," NIST is expressing that all of the five elements are necessary if the term cloud is to be applied:

1. On-demand Self-service
2. Broad Network Access
3. Resource Pooling
4. Rapid Elasticity
5. Measured Service

This definition has been generally adopted throughout the federal government, but it is by no means universal. While NIST states that all five of its essential characteristics should be present for an application to be considered Cloud Computing, other federal officials and experts state that an application that has some, but not all of these characteristics, could still be considered Cloud Computing. [Ibid, GAO Congressional Report above] Whether your organization adopts the NIST definition or creates it own, any definition of Cloud Computing will incorporate many of the characteristics of the NIST definition in [Table 1](#).

Cloud Computing occurs when a Measured Service of a virtually unlimited Resource Pool (i.e., software, platforms, or infrastructure) is accessed over a broad network shared by multiple tenants (i.e., agencies, programs, or applications within an agency) and are rapidly provisioned and released using On-demand Self-service.

TABLE 1: NIST DEFINITION - ESSENTIAL CLOUD COMPUTING ELEMENTS

CLOUD CHARACTERISTIC	EXPLANATION
On-demand Self-service Authorized agencies must be able to provision and release capabilities, as needed, automatically, without requiring human interaction with each services provider.	Usually accessed via a Web browser interface requiring login credentials that permit the acquisition or release of software, platforms, or infrastructure by the agency. If you have ever created an account on eBay, Gmail, Twitter, or GovLoop, then you have used On-demand Self-service.
Broad Network Access Once provisioned, the software, platform, or infrastructure maintained by the cloud provider are available over a network using thin or thick clients.	The NIST definition does not require that capabilities be accessed over the Internet, and does not require that the capabilities be accessible by a Web browser, though these are most common.
Resource Pooling The resources provisioned from the cloud provider are pooled to serve multiple agencies or programs using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to the agency's self-service demand.	This is significantly different from hosting models, where static resources are allocated and permanently assigned to agencies for their exclusive use. Resource Pooling is also why virtualized resources are sometimes erroneously referred to as cloud. Virtualized resources may be pooled in a multi-tenant model and dynamically assigned and re-assigned; however, virtualized resources alone do not necessarily exhibit any of the other essential characteristics necessary to achieve Cloud Computing, nor do they provide any of the flexibility and associated benefits. Simply installing virtual blade servers and storage and transitioning applications to a data center may conserve power, space, and cooling, but it does not implement Cloud Computing.
Rapid Elasticity Cloud Computing capabilities can be rapidly and elastically provisioned and released.	Elasticity is defined as the ability to scale resources both up and down as needed. To the consumer, the cloud appears to be infinite, and the consumer can purchase as much or as little computing power as they need. For example, the July 30, 2009 General Services Administration (GSA) Request for Quotation for Infrastructure as a Service (IaaS) required "provisioning of practically unlimited storage, computing capacity, memory (e.g., at 1,000 times our minimum resource unit metrics)" that can be purchased in any quantity at any time.
Measured Service Cloud resource usage is monitored, controlled, and reported providing transparency for both the provider and consumer of the service.	Measured Service provides the charge-back mechanism to the agency and is the impetus to provision the minimum resources required for use at any time, and to scale resources up or down as demand changes. If the agency or program is paying for a service, this creates pay-as-you-go charging where monthly bills fluctuate based upon the agency's demand. Where no chargeback measure exists, the motivation for customers to release unneeded resources may be greatly diminished.

Cloud Computing Service Models

With the essential characteristics of Cloud Computing defined, it is critical to understand the types of services that are available in a Cloud Computing model. The NIST definition of Cloud Computing defines three service models: Cloud Infrastructure as a Service (IaaS), Cloud Platform as a Service (PaaS), and Cloud Software as a Service (SaaS).


Infrastructure as a Service provides various infrastructure components such as hardware, storage, and other fundamental computing resources. Platform as a Service provides a service that runs over an underlying infrastructure. A platform vendor offers a ready-to-use platform, such as an operating system like Microsoft Windows or Linux, which runs on vendor-provided infrastructure. Customers can build applications on a platform using application development frameworks, middleware capabilities, and functions such as

databases. Software as a Service runs on an underlying platform and infrastructure managed by the vendor and provides a self-contained operating environment used to deliver a complete application such as Web-based e-mail and related management capabilities.

[Figure 1](#) describes the three service models.

The GSA Cloud Computing Technical Framework of [Figure 3](#) mirrors the service models used by the private sector and illustrates that a full spectrum of software, platforms, and infrastructure will be available to federal consumers. At one end of the spectrum, IaaS can provide a virtual machine, operating system, and storage that is maintained by the service provider leaving deployment, configuration, operation, monitoring, and maintenance of the deployment environment, application, and data to the agency. With the IaaS service model, the agency has

FIGURE 1: CLOUD COMPUTING SERVICE MODEL DESCRIPTIONS AND A SAMPLING OF CLOUD PROVIDERS




complete control over the components that it deploys to the cloud, and controls the versions, configuration, look and feel, and functionality of the applications executing on the cloud infrastructure.

At the opposite end of the spectrum, SaaS provides a complete environment where the infrastructure, application, and storage are deployed, configured, operated, and maintained by the service provider. With the SaaS service model, the agency has very little control over the software service being provided, except for those configuration settings that the service provider wishes to expose.

When considering a service model, an agency trades-off the amount of control and customization available to it against the development and sustainment costs for the capability being delivered. IaaS provides extensive management and control of a capability deployed to the cloud while simultaneously leaving the agency with the greatest resource requirements for software licensing, development, deployment, and monitoring.


As illustrated in Figure 2, PaaS provides less control and less resource requirement, and SaaS the least. The NIST service models, along with the GSA Federal Cloud Computing Technical Framework of Figure 3, provide guidance to the agency when considering whether Cloud Computing is appropriate to meet agency software, platform, and infrastructure service requirements.

FIGURE 2: CLOUD SERVICE MODELS



GSA developed the Cloud Computing Technical Framework shown in Figure 3 that maps the classes of information technology typically used by federal agencies to the cloud service models defined by NIST.

FIGURE 3: GSA FEDERAL CLOUD COMPUTING TECHNICAL FRAMEWORK⁴



Cloud Computing Deployment Models

NIST differentiates types of clouds as follows⁵:

Public Cloud

In simple terms, public cloud services are characterized as being available to clients from a third party service provider via the Internet. The term “public” does not always mean free, even though it can be free or fairly inexpensive to use. A public cloud does not mean that a user’s data is publicly visible; public cloud vendors typically provide an access control mechanism for their users. Public clouds provide an elastic, cost effective means to deploy solutions.

Private Cloud

A private cloud offers many of the benefits of a public Cloud Computing environment, such as being elastic and service-based. The difference between a private cloud and a public cloud is that in a private cloud-based

service, data and processes are managed within the organization without the restrictions of network bandwidth, security exposures, and legal requirements that using public cloud services might entail. In addition, private cloud services offer the provider and the user greater control of the cloud infrastructure, improving security and resiliency because user access and the networks used are restricted and designated. A private cloud can be managed by a third party and can be physically located off premises. It is not necessarily managed and hosted by the organization that uses it.

Community Cloud

A community cloud is controlled and used by a group of organizations that have shared interests, such as specific security requirements or a common mission. The members of the community share access to the data and applications in the cloud.

Hybrid Cloud

A hybrid cloud is a combination of a public cloud and private cloud that are interoperable. In this model users typically outsource non-business-critical information and processing to the public cloud, while keeping business-critical services and data in their control. A hybrid cloud can also be instantiated with a minimum level of service provided by a customer's non-cloud infrastructure, with additional capacity provided by an internal or public cloud. Processing requests from the infrastructure "cloudburst" to the cloud.

Note that the cloud deployment models describe the classes of customer for whom a cloud instance is operated, and do not describe who "owns" the cloud or the physical boundaries of the data center from which cloud services are provided.

A selection of service and deployment model are risk-based decisions that define the division of information security roles and responsibilities between an organization and cloud service provider(s). Whatever models are considered, there is a need to define SLAs and to clearly delineate security control responsibilities between providers and customers. NIST is currently working with its members to define interagency security requirements for cloud systems and services and related information security controls from both the moderate and low baselines as specified in NIST SP 800-53 Revision 3.

With the essential characteristics of Cloud Computing, service models, and deployment models defined, we can now discuss some common scenarios that an agency might consider when weighing the performance and resource benefits of Cloud Computing.

Cloud Computing Use Case Scenarios

The use cases described in [Table 2](#) are intended to be typical of agency requirements and are not meant to be a comprehensive list of realizations within a cloud environment. Each use case is described in detail after the table.

Production Surge

Also known as the "cloudburst," this use case augments an application's internal agency resources with cloud resources in response to an increase in demand. When the demand subsides, cloud resources are released. Production surge requirements may be sudden, as in the case of emergency services, or predictable, as in the case of the U.S. Department of Treasury needing to accept dramatically increased numbers of tax returns each April. Production surge using Cloud Computing eliminates the need to build, operate, and maintain capacity to the maximum utilization of the surge. By using computing resources only when required, the cost of computing that is unused during non-surge times is eliminated.

Disaster Recovery

A Disaster Recovery can be defined as a production surge requiring about 75% of the resources of the capability being recovered. In the event of a disaster, Cloud Computing resources are allocated and production processing is transferred to the cloud. Where immediate transition is required, cloud balancing, a technique for routing application requests across applications or workloads that reside in multiple clouds, can be used between private clouds, public clouds, or any combination. Although conceptually simple,

distributing requests in an orderly fashion across multiple clouds requires an intelligent, policy-driven, interpretive methodology to be implemented prior to its execution.

Storage

With private sector storage pricing starting around \$0.15 per GB per month (\$150/TB/month), public Cloud Computing storage can be a cost effective alternative to procuring, operating, and maintaining agency SANs. When public Cloud Computing storage is not an option, creating common virtualized storage across the agency can shorten the acquisition process, as well as reduce both project integration time and the overall cost of sustainment for projects that can use it. In addition, provisioning and releasing storage resources as required tends to reduce the total amount of storage procured; whereas before each project acquired the peak amount of storage that might be required, total agency demand can average peak requirements when the storage demands are not simultaneous.

Development

Agencies often include the specification and acquisition of a project development environment prior to initiation of a project. If rapid ramp-up is required, then agencies might use any available equipment or contract with the private sector to provide a development environment. When this happens, the development hardware and software often does not match the test and production environments, leading to additional rework and schedule and resource impacts. Using PaaS in the cloud, agencies can acquire developer resources as needed and release resources when developers are not using

them in either a public or private cloud. Developers can acquire platforms, develop on them, store their work in cloud storage, release platforms when they go home for the day, and acquire platforms and load their work from storage the next morning. Cloud storage can also be used for version archiving, so that developers can test innovative approaches and concepts and rapidly and completely revert to prior versions if necessary.

Test

Even more so than the development environment, it is sometimes difficult or cost prohibitive to acquire, deploy, and maintain a test environment that is identical to the production environment. These differences can lengthen the transition from test to production and can mask latent defects that seemingly tend to manifest when application demand is most critical. By its very nature, Cloud Computing makes the resources necessary to test and stress applications rapidly available to the test team. If the test and production environments are within the same cloud, the test and production environments are identical and the probability of an undetected defect caused by differences in the environments is greatly reduced. Infrequently run tests such as User Acceptance Test (UAT) or Independent Verification and Validation (IVV) are excellent use cases that can leverage Cloud Computing.

Public Web Services

Federal security standards are most often cited as the most challenging aspect of migrating capabilities to Cloud Computing. When beginning to consider cloud implementations, it is natural to consider first transitioning information which requires the least security protection—generally information that is publicly available. This does not imply that public information does not require protection. Indeed, the unauthorized modification, destruction, denial of service, or redirection of users from government sponsored sites will prove injurious to any agency's reputation. Rather, as with any new model, it is most prudent to start small, pilot results, and select applications with fewer security controls. This will allow the agency to gain familiarity with and develop cloud governance, tools, and techniques and rapidly realize results that can be leveraged into larger and more complex projects.

Although not exhaustive, the use cases presented herein are typical of those found in agencies first approaching Cloud Computing. Each agency should assess whether a fundamental reexamination of their investments in technology infrastructure can be extended to Cloud Computing service deployment models that might be leveraged to provide rapid deployment, scalability, and resource savings.

Security in Cloud Computing

One of the difficulties to discussing security in Cloud Computing is that there is no single "cloud." The sheer number of combinations of the SaaS, PaaS, and IaaS service models and public, private, community, and hybrid deployment models necessitate that cloud security be treated no differently than any agency risk-based approach to an information system. There are however, special security issues that may be introduced by a Cloud Computing environment, including trust, multi-tenancy, encryption, and compliance.

In traditional hosting, each customer provisions their own servers running one instance of an application. In a virtualized cloud environment, one physical server may be used to deliver virtual servers to multiple customers. When an agency uses a public or private cloud, it must be able to establish trust in the provider's security model. Because the agency no longer has physical control of the infrastructure, careful attention should be paid in developing security requirements and periodic auditing. As an additional security measure, encryption of

sensitive data at rest is used to mitigate the risk of unauthorized disclosure. Because most applications running within an agency behind their firewall do not incorporate encryption, modifications are sometimes necessary to meet Certification and Accreditation in the cloud.

Although many security issues are introduced, there are also advantages to having applications operate in a cloud environment. Moving public data to a public cloud reduces the exposure of internal, sensitive data to the Internet. The homogeneity of the cloud environment also simplifies security auditing and testing. An example of leveraging these features is the Federal Risk and Authorization Management Program (FedRAMP), a cross federal government initiative, led by NIST, which creates agreed upon security requirements among federal departments and agencies for certifying and accrediting a cloud-based system. Using FedRAMP eliminates the duplication of effort and cost for each agency working to address security requirements for cloud applications.

The security issues described are solvable using the known processes and techniques for a risk-based approach to information technology security. Industry specific SLAs and security models with cloud providers can also be developed to meet the unique requirements of compliance with FISMA, HIPAA, SOX, PCI, GLBA,

and SAS 70 audit standards. As Cloud Computing continues to mature and agencies gain experience and leverage each other's work, the security challenges described will homogenize and additional opportunities for secure, interoperable, and portable cloud applications will emerge.

TABLE 2: CLOUD COMPUTING USE CASE SCENARIO SUMMARY

USE CASE	EXAMPLE	IMPACT
<p>Production Surge The ability to meet unpredictable surge demands.</p>	<p>Web app usually serves 1,000 customers a day, but must serve 50,000 customers a day during an enrollment period 2 weeks per year.</p>	<ul style="list-style-type: none"> Own servers for typical load, and avoid having 50x capacity unused for 95% of the year. Pay only for the computing resources used. Avoid \$MM license capital expense.
<p>Disaster Recovery Low cost automated disaster recovery.</p>	<p>DR plan requires 75% of production capacity in data centers >100 miles apart. Normally huge duplication on infrastructure and licensing is required.</p>	<ul style="list-style-type: none"> Production capacity can be spread across multiple data centers at normal loads. If one data center is disrupted or incapacitated, the second surges to meet demand ... no duplicate hardware costs.
<p>Storage Low cost storage with the ability to expand on-demand.</p>	<p>Avoid the need to purchase and maintain the latest storage SAN for each program.</p>	<ul style="list-style-type: none"> Provision required storage instantly and pay for what is used. Reduced need to maintain peak storage capacity.
<p>Development A highly flexible and reliable development environment for applications.</p>	<p>Development environments vary greatly and requirements can evolve. Contractors often purchase HW and GFE at great expense to the government and transition becomes cumbersome upon contract closeout.</p>	<ul style="list-style-type: none"> When the government provisions development environments in the cloud, contractors simply access the environments. Change in size or access permission is simple, continuity is easy to maintain.
<p>Test Repeatable on-demand test environments.</p>	<p>Testing activity is typically cyclical for an application in maintenance.</p>	<ul style="list-style-type: none"> Just as people performing tests are cycled into other roles, cloud enables test environments to be expanded and contract with the cycles. Standing up UAT and large performance test environment does not require HW investment.
<p>Public Web Services Distribute publicly available information.</p>	<p>USA.gov, Recovery.gov, U.S. Labor Statistics, DisasterAssistance.gov</p>	<ul style="list-style-type: none"> Provide access to publicly available information without the need to rearchitect data center policies. Have a flexible deployment that can automatically adapt to change in user volume.

Federal Risk and Authorization Management Program (FedRAMP)

FedRAMP recognizes that the individual agency risk management of shared systems leads to duplicative management efforts, incompatible vendor requirements, a slowdown in system acquisition and fielding because of lengthy compliance processes, and has the potential to result in the inconsistent application of federal security requirements. FedRAMP minimizes these by creating a government-wide initiative to provide joint authorization services and continuous monitoring, including unified government-wide risk management. Agencies can leverage FedRAMP authorizations for their Certification and Accreditations while retaining their responsibility and authority.

FedRAMP provides centralized security authorizations of Cloud Computing systems (both commercial and government) to be used government-wide. Centralizing authorizations allows multiple federal agencies to leverage a single security authorization. The FedRAMP Joint Authorization Board (JAB) is comprised of three permanent members – Department of Defense (DoD), Department of Homeland Security (DHS), and the GSA, and the FedRAMP Common Security Requirements (CSRs) are for moderate and low impact systems. CSRs are heightened security requirements from the NIST Special Publication 800-53 baseline that includes 13 additional controls and control enhancements for Low Systems and approximately 50 additional controls and

control enhancements for Moderate Systems. These additional controls address issues of multi-tenancy, shared resource pooling, lack of trust, visibility, and control of the service provider's infrastructure.

FedRAMP makes public a set of CSRs and process documents which the FedRAMP Office and the JAB uses to assess and authorize systems up to FIPS Publication 199 Security Category Moderate and DoD Mission Assurance Category II (MAC II). The DoD will begin transition to FedRAMP in 2011 and will complete the transition by 2014.

Once FedRAMP authorized, it remains the responsibility of each agency to review FedRAMP authorization packages prior to making a decision to accept risk, determine the suitability of the FedRAMP authorization to the agency's mission and risk posture, determine if additional security work is needed, and perform agency specific security activities.

The results are risk management cost savings and increased effectiveness, an interagency vetted approach, rapid acquisition through consolidated risk management, and consistent application of federal security requirements.

Conclusion

Cloud Computing promises to:

- Enable rapid provisioning and deployment of services
- Scale on-demand
- Create services-based environments that are standards-based and interoperable
- Leverage economies of scale
- Promote innovation and service sharing

If successful, the result of transitioning an agency's computing resources to a cloud model promises to significantly reduce the percentage of information technology budget used for operations and maintenance, enabling agencies to reinvest in, and concentrate on, their core mission objectives.

This potential is too great to ignore. As Cloud Computing matures, those agencies that have best developed plans, governance, and pilots exploiting Cloud Computing will be better positioned to rapidly optimize their resources, expand their capabilities, and for some applications, improve the government's ability to create a transparent, open, and participatory government.

Endnotes

1. Peter Mell and Tim Grance, U.S. Department of Commerce National Institute of Standards and Technology (NIST) Computer Security Division, The NIST Definition of Cloud Computing, csrc.nist.gov/groups/SNS/cloud-computing/

2. Federal News Radio, Agencies to justify not using Cloud Computing to OMB, www.federalnewsradio.com/?sid=1836091&nid=35
3. Peter Mell and Tim Grance, page 1
4. U.S. General Services Administration, Federal Cloud Computing Initiative Overview, www.scribd.com/doc/18031511/US-Federal-Cloud-Computing-Initiative-Overview-Presentation-GSA, page 7
5. Peter Mell and Tim Grance, page 2

Other Cloud Computing Sources

- General Services Administration, Federal Cloud Computing Initiative PMO
- Federal CIO Council
- Federal Cloud Computing Initiative Cloud Computing Executive Steering Committee and Cloud Computing Advisory Council
- Apps.gov, www.apps.gov
- TechAmerica Cloud Computing Committee www.techamerica.org/cloud-computing
- Cloud Security Alliance, Security Guidance for Critical Areas of in Cloud Computing, www.cloudsecurityalliance.org/csaguide.pdf
- Cloud Computing Interoperability Forum www.cloudforum.org

Apptis Cloud Computing Leadership

Our contribution to the overall advancement of the federal Cloud Computing direction began in 2008 and continues today through our:

- Participation with NIST in refining the definition of Cloud Computing
- Chairing the TechAmerica Cloud Computing Committed, which seeks to understand and advise both the federal government and Industry on the application and governance of Cloud Computing
- Development of a trusted and highly secure Apptis Cloud Computing Framework
- Partnership with Amazon Web Services™, widely recognized as the cloud industry founder, in addition to other leading cloud providers
- Creation and pilot of the cost effective cloud strategy for FEMA's Disaster Assistance Improvement Program (DAIP), which enabled FEMA to surge to support a 2,500% increase in requests from citizens as part of emergency response
- Working with GSA to provide IaaS on Apps. Gov which will be certified and accredited under FedRAMP, and available for use government-wide

At Apptis, we believe Cloud Computing can help the government invest a greater proportion of its information technology spend on its core missions rather than operating and maintaining information technology infrastructure.

Contact

Mr. Phil Horvitz, *Apptis CTO*

(703) 579-0721
phil.horvitz@apptis.com

About Apptis

Ranked as one of the Top 20 Federal Integrators, Apptis is a leading provider of IT solutions and services for government and industry. Since 1983, we have been applying leading technologies to diverse needs, delivering solutions that are agile, trusted, and business-aligned for optimized performance.

With a fresh-thinking culture steeped in innovation, our goal is to add depth, value, and reach to IT investments by inspiring solutions that set new standards of performance. We do this by leveraging strong domain knowledge and the advanced technical expertise of our people and partners. Integrity and partnership define our approach to client relationships.

Apptis has extensive experience delivering solutions for organizations in key commercial market sectors and all major federal agencies and branches of the military. We employ 1,100 seasoned professionals in 30 states, eight international locations, giving us the professional strength and global reach to deliver solutions anytime, anywhere.

Afterword: To introduce Cloud Computing in the Federal Public Sector, this white paper brings together the work of many public and private sector initiatives. The opinions expressed in this white paper are strictly those of the authors and do not necessarily reflect the opinions of the government or organizations whose work was distilled to create this white paper.

Apptis" is a registered trademark of Apptis Holdings, Inc. All other trademarks, marks, names, or product names referenced in this publication are the property of their respective owners, and Apptis neither endorses nor otherwise sponsors any such products or services referred to herein. © 2010

