SOLUTIONS
WHITE PAPER

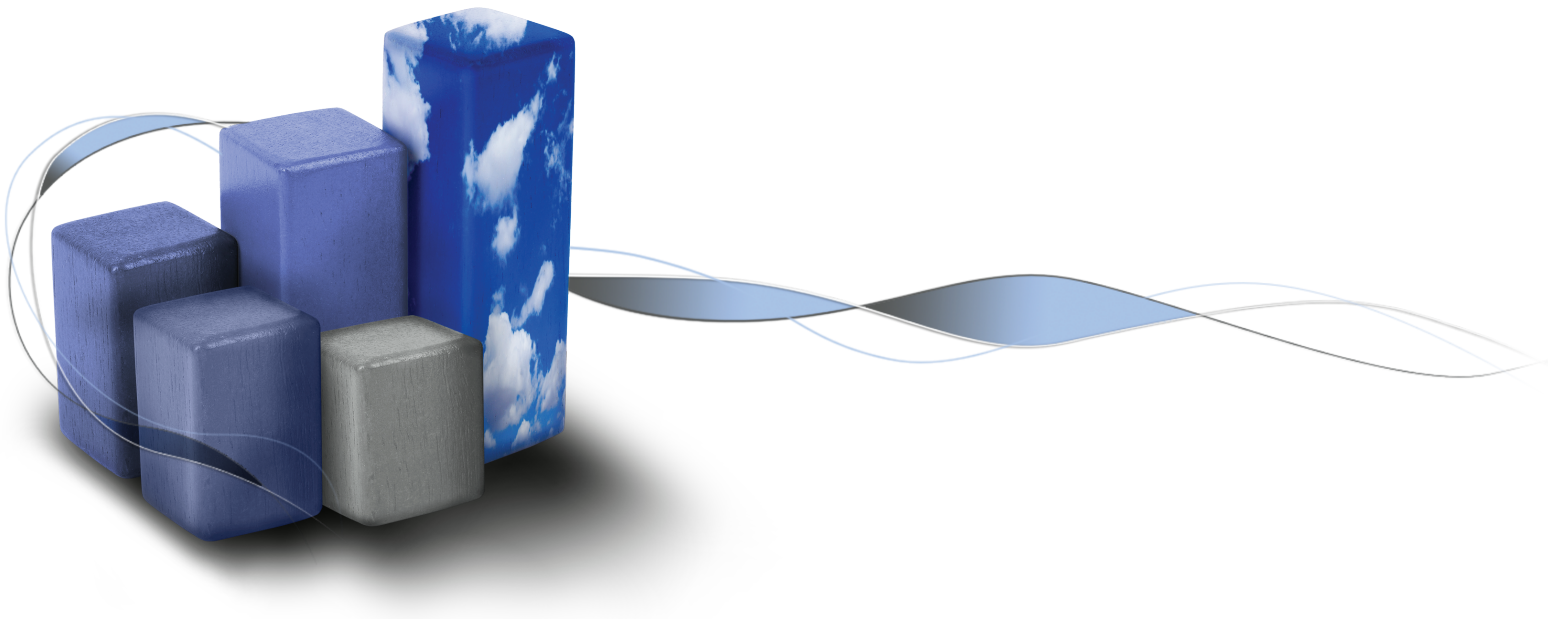

Cloud Computing

Building a Framework for Successful Transition

gtsi®

Contents

Executive Summary

Cloud computing, with the revolutionary promise of computing as a utility, has the potential to transform how IT services are delivered and managed. Yet, despite its great promise, even the most seasoned professionals know little about cloud computing or how to define it. A recent study revealed that 41% of senior IT professionals admit that they "don't know" what cloud computing is. This research follows a similar survey highlighting that two-thirds of senior finance professionals are confused about cloud computing (Version One, 2009).

But as more players enter the market, as proven use drives up the acceptance of the cloud computing model – and budget constraints force IT departments to look for savings wherever possible – the demand for cloud computing solutions is expected to grow exponentially. According to INPUT, estimates are that by FY2013 federal, state, and local government spending on cloud computing will reach $800 million (INPUT, 2009).

The reasons for the increasing interest among government agencies are myriad. To begin, cloud computing offers an entirely new way of looking at IT infrastructure. From a hardware point of view, cloud computing offers seemingly never-ending computing resources available on demand, thereby eliminating the need to budget for hardware that may only be used in high peak timeframes. Cloud computing eliminates an up-front commitment by users, thereby allowing agencies to start small and increase hardware resources only when there is an increase in their needs. Moreover, cloud computing provides the ability to pay for use of computing resources on a short-term basis as needed (e.g., processors by the hour and storage by the day) and release them as needed (Berkeley, 2009).

As for the bottom line, cloud computing enables governments to lower the expense of existing IT services and to cost-effectively introduce enhanced services. Moreover, government agencies not only benefit from increased productivity engendered by cloud computing, but citizens as well benefit from the more efficient use of tax dollars (INPUT, 2009). Costs associated with IT operations in many cases decrease significantly, because services can be purchased on-demand. Finally, administrative time spent attending to the needs of the IT infrastructure can be reduced, with personnel freed to devote more time to an agency's core mission objectives.

"The Federal Government will transform its IT infrastructure by virtualizing data centers, consolidating data centers and operations, and ultimately adopting a cloud computing business model."

(GSA, 2009)

Defining Cloud Computing

The National Institute of Standards and Technology's (NIST) Information Technology Laboratory recognizes that cloud computing is an "evolving paradigm." As such, its definition, attributes, and characteristics are still being debated by the public and private sectors, and are certain to continue to evolve in the near future. Nevertheless, initial steps have been taken toward constructing a universally accepted explanation of cloud computing's key characteristics, as well as definitions for the various deployment and service models. These definitions have been widely reported but are worth repeating, particularly in a field that is still rapidly developing.

According to NIST:
Cloud computing 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.

This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.

The Five Essential Characteristics

On-demand Self Service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.

Broad Network Access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource Pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid Elasticity. Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out, and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured Service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

The Three Service Models

Cloud Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Cloud Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Cloud Computing Framework

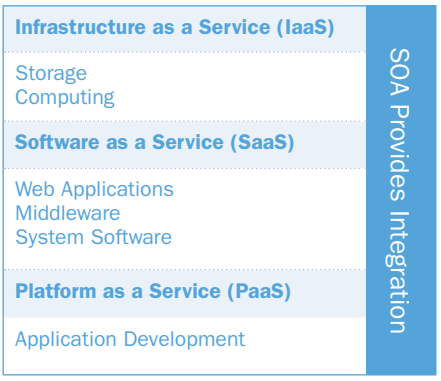

(INPUT, 2009)

Motivators for Implementing a Cloud Computing Solution:

- Increase capacity cost effectively
- Reduce IT operating costs and effort
- Reduce hardware infrastructure costs
- Access SaaS applications
- Free up IT staff

(INPUT, 2009)

Cloud Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

The Four Deployment Models

Private Cloud. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.

Community Cloud. The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.

Public Cloud. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid Cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

The Benefits of Cloud Computing

As cloud computing begins to take hold, several major benefits have become evident:

Costs. The cloud promises to reduce the cost of acquiring, delivering, and maintaining computing power, a benefit of particular importance in times of fiscal uncertainty. By enabling agencies to purchase only the computing services needed, instead of investing in complex and expensive IT infrastructures, agencies can drive down the costs of developing, testing, and maintaining new and existing systems.

Access. The cloud promises universal access to high-powered computing and storage resources for anyone with a network access device. By providing such capabilities, cloud computing helps to facilitate telework initiatives, as well as bolster an agency's continuity of operations (COOP) demands.

Scalability and Capacity. The cloud is an always-on computing resource that enables users to tailor consumption to their specific needs. Infinitely scalable, cloud computing allows IT infrastructures to be expanded efficiently and expediently without the necessity of making major capital investments. Capacity can be added as resources are needed and completed in a very short period of time. Thus, agencies can avoid the latency, expense, and risk of purchasing hardware and software that takes up data center space – and can reduce the traditional time required to scale up an application in support of the mission. Cloud computing allows agencies to easily move in the other direction as well, removing capacity, and thus expenses, as needed.

Resource Maximization. Cloud computing eases the burden on IT resources already stretched thin, particularly important for agencies facing shortages of qualified IT professionals.

Collaboration. The cloud presents an environment where users can develop software-based services that enhances collaboration and fosters greater information sharing, not only within the agency, but also among other government and private entities.

Customization. Cloud computing offers a platform of tremendous potential for creating and amending applications to address a diversity of tasks and challenges. Its inherent agility means that specific processes can be easily altered to meet shifting agency needs, since those processes are typically changeable by making a configuration change, and not by driving redevelopment from the back-end systems (Heyward and Rayport, 2009).

Issues and Risks

One of the key issues in cloud computing is the move towards a multi-sourced

IT environment, where some services are provided in house, some from other government entities, and some from a range of infrastructure, application, and process suppliers in the form of private, public, community, or hybrid clouds.

Any preparation for such transitioning requires a thorough review of an agency's IT strategy in light of its mission needs. Which fundamental capabilities need to stay in house as mission critical? Which are better suited for providers to deliver, and which lend themselves to the pay-per-use cloud approach? These considerations should be made in conjunction with the imperative to consolidate, simplify, and optimize an agency's IT environment, to reduce operational costs and free up investment for other mission-focused initiatives.

Implementing a cloud computing IaaS model incurs different risks than managing a dedicated agency data center. Risks associated with the implementation of such a new service delivery model include policy changes, implementation of dynamic applications, and securing the dynamic environment. Most often, the mitigation plan for these risks depends on assessing the IT services needed to support end users and how they will be delivered, establishing proactive program management, and implementing industry best practices and government policies in the management of that program.

For cloud computing to be widely adopted, assurances must be made that data is not only always accessible, but also totally secure. Agencies will undoubtedly need to actively put in place security measures that will allow dynamic application use and information-sharing to be implemented with the highest degree of security. Indeed, any significant data breach will exacerbate already existing fears about whether data is indeed safe in the cloud.

To enable the cloud and fully realize its potential, certain fundamental elements must be addressed. To begin with, the cloud must function at levels equal to or better than the current IT systems – and must deliver tangible savings and benefits, including raising energy efficiency and reducing environmental impact. Users must be assured of near-ubiquitous and open access via the Internet, and be able to move among the cloud platforms as needed – with the users' rights to the data clearly defined and protected. Above all, as previously stated, user data must be secure at all times.

Cloud Computing Maturity Model

The establishment of a cloud computing maturity model (CCMM) provides a framework for successful implementation. GTSI recommends a phased approach to the CCMM, encompassing five key components:

- Consolidation
- Virtualization
- Automation
- Utility
- Cloud

Step 1: Consolidation

An agency's migration towards cloud computing begins with the consolidation of server, storage, and network resources, which works to reduce redundancy, decrease wasted space, and increase equipment usage, all through the measured planning of both architecture and process.

Consolidation is achieved primarily through virtualization but can also be approached by the use of denser computing hardware or even high performance computing. By boosting the speed of critical processes and enabling greater flexibility, the consolidation of data centers and desktops allows agencies to do more with fewer resources – a significant concern in today's economic environment. Moreover, the shift to a unified fabric provides both physical and virtual access to the storage area network (SAN), creating greater efficiency and cost savings by allowing more storage to be consolidated in the SAN.

Network and application modernization is also an important initial step in enabling the transition to a cloud computing environment. A viable alternative to replacing infrastructure components or rewriting critical applications, modernization promotes communication between older systems and newer solutions, all while preserving the value in existing IT systems. Freed from the bonds of a mainframe environment, critical applications modernized through a service-oriented architecture provide agencies with the increased ability to leverage newer technologies.

As for security concerns surrounding cloud computing, modernization actually works to enhance the security of sensitive information stored on critical applications. When established properly, the cloud platform provides security of all data in motion, traveling between the cloud and the desktop, and all data at rest in cloud storage.

Step 2: Virtualization

Virtualization forms a solid foundation for all cloud architectures. It enables

Cloud Computing Maturity Model

Step 1 Consolidation	Step 2 Virtualization	Step 3 Automation	Step 4 Utility	Step 5 Cloud
Consolidation & Modernization of Resources	Abstraction & Resource Pooling	Adaptive, Secure, & Repeatable	Self-Service & Metering	On-Demand & Scalable
Server Consolidation	Server & Storage Virtualization	Policy-Based Provisioning & Management	Service Metrics & Metering	IaaS, SaaS, PaaS
Tiered Storage Consolidation	Desktop Virtualization	ITIL-Based Repeatable Processes	Service Level Agreements (SLAs)	Service-Oriented Architecture
Consolidation of Network Services	Virtualized Network Services	Multi-Tier Security	Incident Response & Audit	Inter-Cloud Federation
Consolidation of Disparate Applications	Application Virtualization	Multi-Tier Data Recovery	Continuous Availability & Failover	Integration of Web 2.0 & Web Portals
Key Enabling Capabilities				
Consolidation	Virtualization	ITIL Service Management	DR & COOP	Cloud Internetworking
Modernization	Thin Client Computing	Network Security	Risk / Vulnerability Management	Integration
Power & Cooling	Green IT	Data Center Security	Situational Awareness	Provisioning
High Performance Computing	Data Duplication	Infrastructure Protection		

the abstraction and aggregation of all data center resources, thereby creating a unified resource that can be shared by all application loads. Hardware such as servers, storage devices, and other components are treated as a pool of resources rather than a discrete system, thereby allowing the allocation of resources on demand. By decoupling the physical IT infrastructure from the applications and services being hosted, virtualization allows greater efficiency and flexibility, without any effect on system administration productivity or tools and processes.

By separating the workload from the underlying OS and hardware, virtualization allows extreme portability. When extended to every system component, desktop, network, storage, and servers – it enables the mobility of applications and data, not only across servers and storage arrays, but also across data centers and networks. Moreover, through consolidation – one of the critical applications of virtualization – agencies can regain control of their distributed resources by creating shared pools of standardized resources that enable centralized management, speeding up service provisioning and reducing unplanned down time. Ultimately, the result is increased use of assets and simplified lifecycle management through the mobility of applications and data.

Although many agencies turn to virtualization to improve resource usage and decrease both capital and operating costs, the ultimate goal in cloud computing is the use of the abstraction between applications and infrastructure to manage IT as a Service (IaaS) in a true cloud environment.

Step 3: Automation

In this stage, automation optimizes an agency's virtualized IT resources. Through a transformative procedure, the infrastructure is automated, and critical IT processes become more dynamic – and greater control is achieved by trusted policies. With automation, data centers can systematically remove manual labor requirements for run-time operations. Among the various forms of automation in practice today, provisioning automation is perhaps the best known and most often implemented.

Rather than managing underlying infrastructure, agencies in pursuit of cloud computing need to move toward managing service levels based on what is appropriate for the application users, whether it's minimum tolerable application latency or the availability level of an application – whatever are deemed critical factors. In this regard, automation becomes an essential element.

With centralized IT and self-service for end users, automation helps agencies to disentangle themselves from the burden of repetitive management procedures, all while enabling end users to quickly access what they require. Ultimately, automation can help agencies to reduce their operating expenses by:

- Reallocating computing resources on-demand
- Establishing run-time responses to capacity demands
- Automating trouble-ticket responses (or eliminating trouble tickets for most automated response scenarios)
- Integrating system management and measurement

Step 4: Utility

In addition to automation, both self-service and metering – feedback about the cost of the resources allocated – are necessary requirements in creating a cloud service. With breakthrough capabilities for end users and agencies, self-service and metering facilitate not only better IT management but the further extension of the user experience.

In the cloud, there is no intermediary between the user of a resource and the processes for acquiring and allocating resources for critical mission needs and initiatives. Since the user initiates the service requests, IT becomes an on-demand service and the costs of operation drop significantly, because costs are incurred only when the service is used and fewer dollars are spent attending to the needs of the infrastructure.

Essential to IT administration is the question of how to maintain service delivery in a fully virtualized, multi-tenancy environment while at the same time providing the highest levels of security – especially for information and services that might leave the data center. A private cloud utility model answers the question, by enabling agencies to retain the data within their network security while scaling and expanding as user demands change, pooling IT resources in a single operating system or management platform. As a result, anywhere from tens to thousands of applications and services can be supported – and new architectures that target large-scale computing activities easily installed.

Step 5: Cloud

Through cloud internetworking federation, disparate cloud systems can be linked in such a way as to accommodate both the particular nature of cloud computing and the running of IT workloads. This federation allows the sharing of a range of IT resources and capabilities – including capacity, monitoring, and management – and the movement of application loads between clouds. Moreover, since federation can occur across data center and agency boundaries, it enables such processes as unified metering and billing and one-stop self-service provisioning.

With cloud computing, communication increases significantly, as data sharing between previously separate systems is fully enabled – and collaboration within and between government agencies grows exponentially. Ultimately, rather than each agency operating in isolation, constricted by the boundaries of its own data center, not only can services be shared among groups, but also costs can be shared and lessened.

GTSI Delivers on Cloud Computing

Government's current emphasis on optimizing common services and solutions through a cloud computing platform is setting the stage for technological innovation – and giving the government a leading role in the cloud's wider adoption. The vision of cloud computing could radically change the structure of how government entities deliver IT services to their constituents, the data they make available, and the tools they use to collaborate and solve problems.

GTSI closely follows the five-step CCMM described earlier to enable agencies to migrate to a cloud environment. GTSI's cloud computing offering consists of IT infrastructure and services that begin with infrastructure consolidation and virtualization, then span to building the cloud infrastructure and providing secure cloud hosting and lifecycle management services that meet all federal security requirements. GTSI supports the cloud infrastructure including the network, storage, computing, platforms, database and middleware, and cloud applications securely hosted and wrapped around GTSI's unique Technology Lifecycle Management (TLM) methodology.

TLM represents a multi-phased approach that encompasses the planning, design, acquisition, implementation, and management of all elements comprising the IT infrastructure. This model enables agencies to better align their mission objectives with the technology, eliminate the burden and costs of IT asset ownership, and create a more predictable service delivery model for its users – allowing agencies to deliver IT Infrastructure as a Service (IaaS).

GTSI's TLM framework aligns with the GTSI Cloud Computing Maturity Model by delivering:

- Multi-phased approach covering all IT infrastructure elements

GTSI Cloud Computing Services

Assessment	• Infrastructure readiness assessment • Business impact analysis of migration to cloud • High-level solution architecture
Design	• Cloud migration strategy (consolidation, modernization, virtualization, power & cooling, security, compliance, SLA's) • Cloud architecture, design, and adoption roadmap
Deployment	• Implementation • Integration • Pre-production testing & performance tuning
Management	• ITIL-based governance for cloud operations, cloud management, & support • Training & knowledge transfer • Staff augmentation

- Long-term strategic methodology for acquiring, financing, and managing IT
- Alignment of technology and services with business objectives
- Transformation of IT into Infrastructure as a Service (IaaS)

Supporting TLM is a range of GTSI service capabilities, including engineering services, integration services, support services, financial services, program management, e-Business, project management, education, and staff augmentation. GTSI's engineering consulting staff offers years of industry experience in developing and integrating state-of-the-art infrastructure technology. GTSI's seasoned engineers hold hundreds of industry certifications, specializations, and high-level security clearances – the right qualifications needed to plan, design, implement, and operate essential infrastructures for cloud computing. From initial site surveys and cloud assessments to cloud infrastructure design and implementation, GTSI professionals use a definitive Service Engagement Methodology (SEM) as the framework for all solution deployments, which provides essential tools, processes, and procedures for consistent quality delivery.

GTSI's cloud assessment services begin with conducting a site survey and infrastructure readiness assessment and developing a high-level solution architecture. Then GTSI's cloud design services create a cloud migration strategy and cloud architecture. GTSI's cloud deployment services involve implementation and integration services, as well as pre-production testing and performance tuning. Meanwhile, GTSI's cloud management services offers ITIL-based cloud operations, management, and support, as well as staff augmentation.

Transforming Government

GTSI's methodology has been applied to a number of federal, state, and local government projects related to infrastructure modernization, consolidation, virtualization, disaster recovery, and staff augmentation – the building blocks to cloud computing. Whether it is server and storage consolidation for NIH, thin client (diskless) computing for the Department of Energy, server virtualization for the DHS Customs and Border Protection, or data center modernization for Army LOGSA, these essential steps toward the cloud have the potential to radically transform how government does business, both internally and externally.

Conclusion

With its convenient, on-demand model for network access to a shared pool of configurable computing resources, cloud computing is rapidly emerging as a viable alternative to traditional approaches – and is carrying a host of proven benefits to government agencies. Costs are being significantly reduced, along with personnel time spent on computing issues. Storage availability increases, high automation eliminates worries about keeping applications up to date, and flexibility and mobility are heightened, allowing workers to access information anytime, anywhere.

Cloud computing can be rapidly provisioned and released with minimal management effort or service provider interaction. Ultimately, with its offering of scalable, real-time, internet-based information technology services and resources, the cloud can satisfy the computing needs of a universe of users, without the users incurring the costs of maintaining the underlying infrastructure.

For more information on GTSI and our cloud computing offering, visit GTSI.com or call us at 800.999.GTSI.

References

GSA, *U.S. Federal Cloud Computing Initiative Overview Presentation,* Scribd.com, 2009

Hewlett Packard Laboratories, *Taking Account of Privacy when Designing Cloud Computing Services,* 2009

INPUT, Federal Industry Insights, *Evolution of the Cloud: The Future of Cloud Computing in Government,* March 2009

Intel, *Developing an Enterprise Cloud Computing Strategy,* 2009

National Institute for Standards and Technology, *Working Definition of Cloud Computing,* 2009

Office of Management and Budget, *Optimizing Common Services and Solutions through a Cloud Computing Platform,* 2009

Rayport, J. and Heywad, A., *Envisioning the Cloud: The Next Computing Paradigm,* 2009

Sun Microsystems, *Introduction to Cloud Computing Architecture,* 2009

University of California at Berkeley, *Above the Clouds: A Berkeley View of Cloud Computing,* 2009

Version One, *Cloud Confusion amongst IT Professionals,* 2009.

Supplemental Material

Intel, *Developing an Enterprise Cloud Computing Strategy*

Hewlett Packard Laboratories, *Taking Account of Privacy when Designing Cloud Computing Services*

Mary Johnston Turner, Frank Gens, IDC, *Cloud Computing Drives Breakthrough Improvements in IT Service Delivery, Speed, and Costs*

NetApp, *NetApp Architectural Plans for Enabling a Dynamic Data Center for Today and In the Future*

Sun Microsystems, *Introduction to Cloud Computing Architecture*

VMware: *Clearing the Fog for a Look into the Clouds,* Mark Bowker, Enterprise Strategy Group

About the Authors

Prem Jadhwani
Senior Solutions Architect, GTSI

Prem Jadhwani has 12 years of experience working in the enterprise IT space with both federal and commercial customers. Currently at GTSI, Mr. Jadhwani defines the solutions, capabilities, standard operating procedures, and “go-to-market” strategy for solutions involving networking, security, and emerging technologies. Mr. Jadhwani holds a MS degree in computer science and a MBA degree in marketing and strategy from the Illinois Institute of Technology in Chicago. He has completed his graduate work towards a Ph.D. degree in Information Security at George Mason University in Virginia and holds a number of professional certifications including CISSP, CSSLP , CISM, and GSLC.

About GTSI

GTSI Corp. is the first information technology solutions provider offering a Technology Lifecycle Management (TLM) approach to IT infrastructure solutions delivered through industry-leading professional and financial services. GTSI employs a proactive, strategic methodology that streamlines technology lifecycle management, from initial assessment to acquisition, implementation, refresh, and disposal. TLM allows government agencies to implement solutions of national and local significance quickly and cost-effectively. GTSI's certified engineers and project managers leverage strategic partnerships with technology innovators. These experts use proven, repeatable processes to design, deploy, manage, and support simple to complex solutions, to meet governments' current and future requirements and business objectives. GTSI is headquartered in Northern Virginia, outside of Washington, D.C.

2553 Dulles View Drive, Suite 100, Herndon, VA 20171-5219 · 800.999.GTSI · GTSI.com