State of Ohio, Office of Information Technology

Cloud Computing Guidelines

This document is intended to highlight the State's position with respect to the appropriateness of each cloud service offering, identify key considerations that may factor into decision making with respect to solution selection, proposed investments, guidance to Agencies and vendors as well as to articulate the anticipated implementation horizon for the State's private cloud in the context of emerging public cloud offerings.

Current Situation

The State is operating a highly complex and distributed IT infrastructure that, via more than 30 data centers, supports more than 1,600 applications via 5,500 servers and hundreds of private and public network elements. The State wishes to dramatically simplify these applications and the supporting infrastructure while reducing operating and future costs associated with providing IT Infrastructure to State Agencies.

Central to this effort will be the expansion of, and Agency migration to utilizing State private (and incorporation of public) cloud computing which is designed to provide the State a secure, high-performance and dependable foundation for computing at a cost point that between 50-75% less than services offered today. In total the State spends approximately $108M annually for IT infrastructure and related services, hardware and labor and the goal of the State's cloud effort is to substantially reduce spend and reapportion that spend into applications and services designed to support the Citizens and Businesses of the State of Ohio.

In fiscal year 2010 the Office of the State CIO in concert with the State's CIO Leadership Management Council (LMC) developed a Statement of Direction with respect to reducing and realigning the cost, methods and strategies associated with delivering IT services to the State. In fiscal year 2011, the State began developing implementation phasing strategies for the Statement of Direction as a whole with a specific emphasis on IT Infrastructure consolidation as a logical (and significant) first step in realizing the Statement of Direction. In parallel, a position paper with respect to cloud computing services has been drafted that describes the State's position with respect to private and public cloud models and specific service models such as cloud-based: Software as a Service (SaaS) , Platform as a Service (PaaS) and Infrastructure as a Service (IaaS).

The State has identified several public cloud solutions available on the marketplace. Based on the relative merits of these offerings believes a clarification with respect to applicability, usage and other factors is an essential element of the State's strategy to implement cloud services. Current market leaders were selected for an evaluation of relative features and capabilities of their platforms in comparison to existing State capabilities and planned cloud services.

The selected vendors in this document should in no way be construed as a validation or endorsement on behalf of the State of these vendors, nor should any party imply that the State has a preference, plan or predisposition to utilize these (or other) vendors for the provision of public cloud services. They were selected purely based on visible market share and the availability of public data and analyses of their services.

Fit/Gap Assessment of Current Leading Marketplace Offerings

In consideration of the State's application and IT Infrastructure that supports these applications and based on ongoing project demands of large Agencies summary requirements for the Infrastructure (e.g., facility, servers,

storage, network connectivity, backup etc.) as well as IT application profile (e.g., scope, business process support, availability, redundancy/reliability, security considerations etc.) the State's use of technology is both highly complex and highly fragmented.

Notwithstanding financial considerations addressed elsewhere in this document, a technical and functional fit/gap has been constructed. For this example: the State Cloud refers to existing and planned OIT State Private Cloud offerings; Hybrid refers to either dedicated hosted SaaS solutions from an application vendor as an integrated solution or a SaaS offering that maintains data with a secure State facility, but leverages software and other solution elements via a cloud offering (e.g., hosted Microsoft Office applications); and the Public Cloud refers to a composite of leading public cloud offerings such as Microsoft Azure and Amazon EC2.

Requirements and considerations in this section are presented in summary form to illustrate key functional, technical and operational differences between each cloud offering and are meant to be representative as opposed to complete. Structured sourcing efforts would traditionally include significantly higher degrees of detail and rigor that, in general, may pose additional challenges for public and hybrid cloud solutions.

Legend: ✓Well Suited, △Proceed with Caution, ⃠Not Well Suited

Requirement Area	Key Considerations	State Cloud	Hybrid Offering	Public Cloud
Infrastructure Requirement: Attribute Summary				
Security and Privacy	▪ Maintenance of personal, private and sensitive data ▪ Resiliency to unauthorized access	✓	⃠	⃠
Technical Performance	▪ High CPU, Memory, Bandwidth or I/O Requirements ▪ Predictable workloads	✓	✓	△
Availability & Service Levels	▪ 24x365 availability, 99.95%+ uptime ▪ Fault tolerance, redundancy	✓	△	⃠
Customization	▪ Standards enforcement (OS, DBMS, Security, System Image) ▪ Tailored to Application / Agency technical requirements within standards	✓	△	⃠
Cost Savings Impact Areas	▪ Operational Cost of Ownership ▪ Ongoing TCO reduction, Cost avoidance	✓	✓	✓
Driver of Statewide Consolidation	▪ Reduction in systems, software and application counts, operational complexity ▪ Simplification of integration, workflows and labor requirements	✓	✓	✓
Migration Profile	▪ Ease of migration from current solution platform to cloud based offering ▪ Technical migration complexity profile	△	△	△
Integration (Process & Technical)	▪ Cross system workflow support and data exchange ▪ Mixture of sensitive and non-sensitive data ▪ Adherence to State integration standards	✓	△	⃠
IT Application Profile: Suitability to Task				
Websites and Public Interaction (Informational)	▪ Presentation of State / Agency presence to public / businesses ▪ Distribution of non-sensitive data	✓	✓	✓
Transactional Websites	▪ Collection of non-sensitive transactional data ▪ Collection of low-risk fees/revenue or other information	✓	✓	△
End-User Computing	▪ Common / routine office productivity tasks ▪ Workflow data sharing, reference data, non-transactional data hosting	✓	✓	✓
Workgroup Enablement	▪ Storage of routine forms, data, knowledge management and other workgroup enablement data / functions	✓	✓	✓
Business Process Enablement	▪ Integrated processes within a single application or application suite ▪ Processing of transactional data non-critical to the State or public safety, revenue collection	✓	△	△
Standalone Operational Systems	▪ Agency specific and non-critical applications ▪ Simple integration and reporting ▪ Routine Agency functions (non-sensitive data)	✓	△	⃠
Cross-Agency Systems	▪ Agency specific critical applications ▪ Complex integration and reporting ▪ Routine Agency functions (sensitive data)	✓	⃠	⃠
DR – Non Critical Systems / Data	▪ Data replication of non-sensitive systems and data ▪ Archive and reference data management	✓	✓	△
State ERP (OAKS)	▪ Operational Uptime and Performance	✓	⃠	⃠

Requirement Area	Key Considerations	State Cloud	Hybrid Offering	Public Cloud
	▪ Highly complex business rules and integration ▪ Maintenance of Sensitive Data			
Highly Integrated Operational Systems	▪ Complex integration and workflows, potentially spanning many systems and work groups ▪ High operational uptime and performance requirements ▪ Maintain personal or confidential data	✓	⊘	⊘
State Critical Systems	▪ Systems that directly influence the State's ability to perform Public Safety, Citizen Services, Revenue Collection and/or Critical Employee Services	✓	⊘	⊘

Financial Comparison

A summary budgetary comparison between status-quo Statewide Agency was constructed utilizing:

- FY11 Actual State Infrastructure Costs inclusive of data center facilities, infrastructure hardware and software, network connectivity (but not bandwidth) and State labor;
- Published FY2012 OIT Rates for Hosted and Managed Servers;
- A pricing sample from the private sector utilizing a leading strategic outsourcing vendor;
- Published Amazon EC2 rates from http://aws.amazon.com/ec2/pricing/ as of August 2011; and
- Published Microsoft Azure pricing from http://www.microsoft.com/windowsazure/pricing/ as of August 2011.

For this analysis it was assumed that all 5,500+ servers would participate in a cloud service of some sort, true 24x7 processing and availability. Due to the nature of Statewide IT Infrastructure computing, disk usage and transactional I/O rates were not available and therefore were not factored, but in any case would be additive to the Amazon and Microsoft pricing based on actual usage.

From a commercial contracting perspective, published pricing does not factor a competitive procurement, related negotiations and other factors which may offset some of the disk storage, networking and transactional I/O fees associated with a true procurement.

It should be noted that approximately 30% (or 1,650) State servers would be classified as meeting a true 7x24 application availability window hosting approximately 490 State critical applications. As the State does not maintain centralized criticality profiles (e.g., 8x5 business hour, 12x7 extended or 24x7 continuous use) for its application inventory, applying actual use profiles to each application may result in a reduced processing requirement and therefore cost to the State when considering Amazon and Microsoft pricing.

Scope / Attribute	State Status Quo (Estimate)	State Cloud Services	Outsourcing Vendor	Amazon EC2	Microsoft Azure
Hosted "XL" Server (Annual Cost)	$20,596	$2,772 hosted $4,932 managed	$3,566	$7,100 (1,2)	$8,415 (1,3)
Pro-forma pricing, addressable State Servers (5,500)	**$108.75M**	**$25.33M**	**$18.32M**	**$39.05M** (1,2)	**$46.28M** (1,2)
Comprehensive SLAs	Best Effort	Yes	Yes	No	No
Break/Fix	Yes	Yes	Yes	Partial	Partial
Patch Management	Partial	Yes	Yes	No	No
Virus / Intrusion Detection	Inconsistent	Yes	Yes	Partial	Partial, Additional Price
Transactional I/O	Unlimited	Unlimited	Unlimited	Additional Price	Additional Price
Technology Refresh	Budget Permitting	Yes	Yes	Additional Price	Additional Price
Bandwidth	Included, significant redundancies	Included	Included	Additional Variable Cost	Additional Variable Cost

Notes:

(1) Plus bandwidth / connectivity charges

(2) Plus I/O and Disk Storage beyond 15GB

(3) Plus database charges beyond 5GB

Based on this analysis and the available data, a wholesale migration of State functions to public cloud providers may not be prudent at this time from a financial perspective. However, given the State's application portfolio and in consideration of several large Agency systems development lifecycle (SDLC) projects, private cloud offerings may have a positive financial outcome under one or more of the following business scenarios:

- Seasonal usage of systems in support of non-routine processes;
- Leverage of commercial off the shelf software implemented in the cloud;
- Archive or backup of non-sensitive data, images or information;
- Support of project activities and environments such as development, prototyping, testing, configuration management, training, knowledge and project management;
- Standalone systems that support smaller workgroups or highly structured tasks;
- Non-secure transaction processing or trivial information exchange (e.g., distributing forms or bulletin);
- Maintaining reference data or standalone databases with predictable transaction profiles; and
- Other applications that require business (or extended) hour support only on a business case basis.

However, in the absence of detailed server usage profiles (i.e., cpu, memory, storage etc.) it's a mistake to look at cloud pricing and conclude that public cloud offerings are more expensive. Public cloud providers, under the appropriate conditions, offer other cost advantages:

Pricing transparency will allow an Agency to determine how much it will cost to run a server per hour. The pricing is available online, so in general there are no surprises other than the noted disk and bandwidth utilization. Given the flexible provisioning and pricing tools, should an Agency decide to alter configuration details, the State could readily calculate what the new pricing will be, unlike strategic outsourcing contracts where change orders are seen by outsourcers as a margin bump opportunity or fall into relatively complex ARC/RRC (additive/reduced recurring charge) calculations.

Public cloud pricing is fixed you can calculate what the total cost per month will be. The State should not underestimate the attractiveness of certainty.

Purchasing and provisioning is relatively easy: in general need nothing more than a credit card to get going with major providers which in some cases would allow the State to potentially avoid the need to endure visits by sales people, create RFPs, buy machines, process invoices etc. The period from decision to do something to actually getting to work is minimized, enabling organizational agility at the possible expense of cost containment and standards control.

As a side note regarding power cost and availability, in general the cost just to power a server is not much less than $0.10/hr (or approximately $876 a year); when factoring the total cost of the machine itself, procurement, maintenance, updates and other costs, and public cloud services may well be cheaper on a per-hour basis for periodic use situations.

Public Cloud Robustness

In consideration of the current state of the art in the provision and operation of public cloud services and based on recent outages it is also clear that this initial Government implementation may not be ready for prime time, especially for systems or applications requiring a redundant architecture.

Amazon EC2 runs in a number of locations (known as regions) across the United States (east and west). Govcloud® is located in the west region which has some geographic diversity appeal. In April, 2011, Amazon Web Services

(AWS) had a major fault which initially impacted multiple availability zones in one region. After approximately 3 hours, Amazon was able to confine the outage to one availability zone.

Users that were implemented using those sites by Amazon that did not include provision for redundancy or failover (by having servers running in multiple regions and load balancing between them) were out for a significant amount of time (in some cases for days). It is unclear as to whether this was a customer preference or a design consideration of the Amazon service.

Sites that primarily use Amazon Web Service (AWS, a complimentary cloud service) for their processing, had architected their systems to use multiple regions (e.g., Netflix), had localized outages, but were able to restore services fairly quickly. At this time AWS government cloud does not offer regional redundancy. GovCloud® is one region with multiple availability zones. Government customers can architect a system to use multiple availability zones within one region (the GovCloud® region) but not across multiple regions for redundancy because the appropriate government security requirements are not currently implemented in the other regions. If the west data center goes down (or the GovCloud region goes down), any systems or applications may be adversely impacted until Amazon is able to rectify the problem. It is clear at this time as to Amazon's plans to fix this issue in the long run, but for now, it is not an appropriate solution for the State.

Similar availability, reliability and architectural analysis are ***strongly advised*** for the State prior to contemplating any public cloud offering. Points of failure, redundancy and vendor track record are essential minimal pre-requisites to utilizing any public cloud offering.

Implications to State's Cloud Strategy

In consideration of the State's 1,620+ applications, more than 5,500 servers, and 4+ petabytes of data stores and in light of the State's 24x7 public remit for many services to the citizens and companies doing business in the State of Ohio, an expansion of centralized State cloud and leverage of public cloud services must be implemented in a controlled and balanced manner that factors complexity, risk, financial rewards and security and privacy considerations.

However, it seems clear at this time that in the foreseeable future, several public cloud options may be viable to help the State address non-critical, seasonal or sporadic use applications as well as non-production environments such as development, testing, training and demonstration instances.

Additionally, in consideration of the relative age of several of the State's applications and the move of the industry to offering Software as a Service (SaaS) offerings, the migration to public cloud service replacements for these applications may be unavoidable. Therefore for these types of offerings the State must adapt its IT solution procurement activities to accommodate SaaS and factor not only solution functionality, but operational, service level, maintenance and other considerations that come hand-in-hand with SaaS.

As a general strategy the State's cloud strategy could be viewed simplistically as:

State Application Requirements (Functional, Integration, Technical, Security etc.)		
Private State Cloud / State Hosting	Hybrid Private / Public Cloud	Public Cloud
▪ Public Safety or Life Critical Services ▪ Essential Citizen Services ▪ Legal / Regulatory / Statutory Compliance ▪ Revenue Collection ▪ Critical Employee Services ▪ SLA critical solutions (24x7, 99.9%+ uptime) ▪ Complex, multi-system workflows and integrations	▪ New Agency SaaS systems ▪ Legacy replacement SaaS offerings ▪ Workflow / Imaging Systems (data maintained in State) ▪ Integration with Federal and Local Government Cloud(s) ▪ Transactional Systems maintaining non-sensitive Data ▪ eMail (once standardized and	▪ Common End-User Computing and Productivity Software (e.g., MS-Office, common desktop tools) ▪ Workgroup Productivity Applications ▪ Agency Web Presence (non-Secure) and Information Distribution ▪ Non-sensitive data storage, reference data storage, archive ▪ New Agency SaaS systems (non-sensitive

State Application Requirements (Functional, Integration, Technical, Security etc.)		
Private State Cloud / State Hosting	Hybrid Private / Public Cloud	Public Cloud
▪ Personal / Private / Sensitive Data	consolidated) ▪ On-demand SDLC Project Environments ▪ DR of non-critical systems, non-sensitive data replication ▪ Infrequent use applications, simple cross-system workflows and integrations ▪ Non-SLA sensitive offerings (best effort, reasonable extended hour support)	data) ▪ Legacy replacement SaaS offerings (non-sensitive data) ▪ Infrequent use applications, potentially standalone workflows and integrations ▪ Non-production, SDLC environment hosting (no sensitive data) ▪ Non-SLA driven offerings (best effort, reasonable business hour support)

Cloud Implementation Opportunity

For the past biennium, State Agency IT organizations have been working to move to contemporary computing platforms that are virtualized under a set of agreed standards and operating practices. As of the end of FY11 some 3,601 physical servers were identified in the State, and 1,357 virtual machines were implemented. Based on a survey of these servers, OIT determined that, in aggregate, cpu utilization on average for these servers was 4.48% and measured peak utilization was approximately 8.8% which suggests that there remains a significant opportunity to the State to better utilize computing assets through continued virtualization. Assuming a plausible scenario using these statistics, and in consideration of the State's 5,500+ servers (both surveyed and surveyed) it is not inconceivable for the State to require between 600 (theoretical limit) and 1,000-1,200 server images (practical reality) to conduct State business.

A wholesale shift of the status-quo environment without any optimization (e.g., application retirement and elimination, storage rationalization, network simplification, and more aggressive virtualization) is not advised. The migration of a physical server to a public cloud without this optimization essentially reduces costs based on the economics mentioned elsewhere in this document but does not address the large scale server reduction (5 to 10:1) that will drive significantly increase savings to the State while reducing operating complexity and capital demands.

In short, the easiest application to "migrate" is a legacy application that is no longer required by the State (i.e., retire the software and server asset and avoid migration costs). In addition to applications that generally fit the "public cloud" profile in the prior section, applications that are already running on virtual machine instances would be well suited to migration as OIT, most leading outsource and public cloud providers offer "virtual image import" capabilities. Applications that are as yet not virtualized (but scheduled) would be more complex, followed by applications that due to technical or operating considerations defy virtualization and/or applications that maintain private or sensitive data.

To provide the State as an "order of magnitude" estimate based on FY2010 and 2011 application inventories and FY2011 server inventories the following table is provided.

Cloud Target	Private State Cloud or Dedicated State Hosted Services	Hybrid Private / Public Cloud	Public Cloud
Agency Applications (Total)	973 Agency Specific Applications 259 Financial or HR Applications ➔ 68 interagency systems ➔ Retire obsolete applications Enforce Standards	←212 Common Applications➔ *some financial applications* *some interagency systems* Retire obsolete applications Enforce Standards	*some common applications* 90 Public Interactive Websites 26 OIT Consolidated eMail and Collaboration systems Retire obsolete applications Enforce Standards
Agency Servers	4,500+ virtualized to <800 Retire obsolete servers	750+ virtualized to < 150 Retire obsolete servers	300+ virtualized to < 75 Retire obsolete servers
Storage	*Data not available, anticipated to be similar to the above*		
Local and Wide Area Networking			

Current Environment Implications

In consideration of the State's current highly fragmented and distributed infrastructure, moving forward to implementation of these cloud guidelines requires several activities that need to be well planned and executed. In summary view, these activities are arranged in the following four areas: technology, people, communications/change management, and governance. A summary view of each with brief commentary follows:

Technology

The technology thread sets the foundation for the consolidation of IT services Statewide. Central to the full deployment of the State's private cloud, and by extension leverage of public cloud services is the remediation of the SOCC and offering centralized IT Infrastructure services Statewide. This step is designed to:

- substantially eliminate duplicative data processing facilities;
- provide alternate processing sites for business resilience and disaster recovery services;
- materially reducer server and storage foot prints through virtualization and standardization;
- establish the basis for telecommunications networking consolidation; and
- deliver a reliable computing infrastructure foundation for applications and IT services Statewide.

In concert with these infrastructure activities, the State should identify superfluous, redundant or other applications that can be retired.

People

People are an essential step in the realignment of IT costs is moving infrastructure staff to a central organization to provide IT Infrastructure services to Statewide Agencies in support of application and service offerings to the constituents of the State. As part of this activity, the State should evaluate the IT labor pool in light of:

- anticipated or planned retirements of the workforce, and skills required to operate a consolidated cloud infrastructure cost effectively and reliably;
- the required staffing levels associated with the operation and maintenance of a highly virtualized and homogeneous server, storage and network offering realizing that the operation of 750 to 1,000 server images from a centralized body requires a substantially smaller labor pool;
- clear definitions of roles and responsibilities of the IT Infrastructure Services organization with respect to operational and maintenance activities in support of Agency application development and maintenance activities;
- a required change in orientation and culture to function as a high performance service organization that is measured (and potentially compensated) based on achieving operational and financial excellence; and
- equipping the workforce with the requisite tools, training and opportunities to continue to advance the state of the art with respect to IT Infrastructure services.

Communications & Change Management

As the State has historically operated in a highly fragmented and federated model, the change to a central services provider model enabled by the cloud will undoubtedly be profound. It is essential that Agencies are motivated to participate in the program actively and, if called upon, submit their "best and brightest" infrastructure architectures, designs, learning's and importantly personnel to supporting this effort. Central to affecting this change is:

- a concerted communications and participation program that is designed to ensure understanding and drive aggressive Agency support and adoption of this program;
- active participation of Agencies and identification of strategies to drive early, visible and measureable successes as a result of moving from Agency-specific infrastructure to the State's IT Cloud;
- pragmatic planning processes that balance financial realities, business risk and financial rewards but are honest enough to address the "newness" and "different than usual" aspects of delivering IT Infrastructure services via the cloud;
- development of repeatable methods to migrate, virtualize, test and transfer services from an Agency to the central services organization – initially within the SOCC, then onwards to other Statewide data centers; and
- regular updates of successes, learning's and enhancements to impacted Agency Stakeholders.

Governance

Governance is key to moving the State forward. It is important that Agencies understand and are incented to comply with this initiative and actively participate without exception. It is equally important that IT Infrastructure consumers (Agencies) have a say and view into the service as well as:

- motivation to participate and not continue to add to the complexity of the program by continuing to invest in and deploy IT Infrastructure elements;
- oversee program execution and work collaboratively to adjust course and plans if and when required;
- understand the underlying economics (both cost and savings) through the State's implementation of this program and provide support to complete the program successfully;
- see the true "value for rate" in the form of increased service levels, reduced capital and operating expenses and work to continue to drive the overall cost of Information Technology within the State down; and
- build on early successes to drive Infrastructure consolidation forward to the greatest effect and unlock network, application and services consolidation as a result of effective governance and stewardship of the program.