

# How to migrate your business to the Cloud



## Introduction

As discussed in Part 1 of this Cloud migration white paper series – “Should I migrate my business to the Cloud” – there are many benefits to using Cloud based services to host key business systems and IT infrastructure. They center on higher uptime, improved redundancy, and latest feature-set benefits as compared to typical in-house IT deployment.

The previous white paper explained these benefits and highlighted how to audit your current IT environment to prepare for such a transition. Part 2 of this white paper series details different scenarios for Cloud migration tailored to differing levels of in-house expertise, business continuity goals, and preferred scope of Cloud adoption.

The diagram below illustrates three scenarios of Cloud migration on a scale of increasing complexity. The description for each scenario follows with the objective of helping you decide which is best suited for your organization.

EASY		MEDIUM	COMPLEX
Self Contained Applications	Migrate your Server Room to the “Cloud”	Build a hybrid “Cloud” and on-premise environment	

## EASY ►

### Start with simple self-contained applications.

The simplest way to get started with the Cloud is to subscribe to self-contained Software-as-a-Service offerings. These Cloud services are considered “Easy” because they can be kicked-off with limited in-house IT expertise. Often the Cloud service will include on-boarding support, self-service migration tools, and FAQ type documentation to ease transition to the new services.

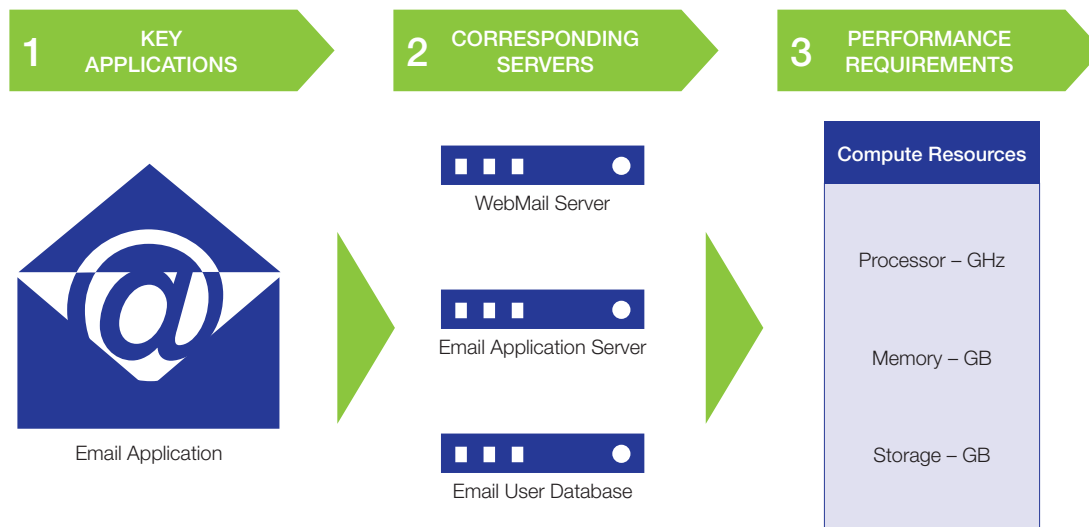
Typical examples of self-contained SAAS offerings include CRM services (Salesforce.com), hosted email services (Office 365), or hosted Unified Communications offerings (TELUS Cloud Collaboration Services). In most cases, these applications are implemented by making simple configuration changes (eg: Domain Name Service (DNS) change) to re-direct your employee traffic from the in-house server to the newly subscribed Cloud based offering. The in-house servers can then be de-commissioned without affecting other applications still hosted within the corporate server room.

In some cases, the Cloud services – eg: Hosted Contact Centre – are designed to easily integrate with existing on-premise equipment – eg: on-site telephone system. In this example, contact centre agents can use any phone – office, home, or cell – to take customer calls while accessing leading edge contact centre capabilities – customer records (CRM), skills based call routing, web chat – through any Internet feed. Businesses can quickly improve customer service through the use of these capabilities as well as business continuity planning – via remote-agent strategies – all without having to forklift out the current telephone system.

# MEDIUM Complexity ► Migrate your Server Room to the Cloud.

As previously mentioned, a key driver for migrating mission critical IT infrastructure to the Cloud is to circumvent the short comings of an outage prone on-premise server room. This Cloud migration strategy is deemed to have Medium complexity because extensive due-diligence and planning will be required before migrating the corporate systems to a new Cloud environment. Often organizations hire IT consultants to help perform this transition. Every organization has different infrastructure and requirements so self-service tools or automated onboarding is no longer an option.

Pre-migration activities include an application and IT infrastructure audit to compile a comprehensive list of all systems and hardware needed in the new Cloud environment. It is important to map all linkages between different application servers so they are migrated as a cohesive unit. Often an application will be composed of multiple servers – front end, application, and database servers – all of which must be migrated together to keep the system running. Once you have a list of (1) all the key applications, and (2) the corresponding servers from which they are composed, you are ready to determine (3) the performance requirements for your future Cloud environment – see diagram below.



Cloud based Infrastructure as a Service (IaaS) offerings are typically sold on a Virtual Machine (VM) basis. Clients subscribe to multiple VM hosted in the Cloud; each representing a server that was previously hosted on-premise. Per-VM pricing varies depending on the amount of computing resources (CPU Processing, Memory, and Storage) dedicated to each VM. Other offerings (eg: TELUS Flex Cloud) allow you to subscribe to a pool of resources, granting you complete flexibility to turn-up and down VMs in real-time, oversubscribe the environment using VM priority levels, and piggy back temporary test-and-dev servers on top of your production resource pool.

Regardless of the IaaS subscription model, a network audit is needed to understand current compute resource utilization levels for each of the servers in your environment. This information is critical to ensure you don't under provision (or over provision) your servers in the new Cloud environment.

A Virtualization Assessment will study the compute resource consumption patterns of your server environment over a period of time – often two weeks. Based on the results you will know on a per server basis the corresponding CPU, RAM, Storage, and Network subscriptions you will need for the corresponding Cloud VM images. Virtualization Assessments can be performed by the Advisory services division of your service provider, or done in-house using freely available tools including:

- Microsoft Performance Monitoring (PERFMON):  
<http://technet.microsoft.com/en-us/library/cc749249.aspx>
- Microsoft Assessment and Planning Tool (MAP):  
<http://www.microsoft.com/en-ca/download/details.aspx?id=7826>
- VMWARE Capacity Planner for Server Consolidation:  
<http://www.vmware.com/products/capacity-planner>

Once you know how many VM to set up in the Cloud environment, and the corresponding performance requirements, you are then ready to plan the migration process. If you already have a Virtualized environment, you can start uploading the existing VM images to the new Cloud environment. Typically the necessary VM files are bundled into an open standards based package: Open Virtualization Appliance (OVA). Your in-house virtualization platform – typically VMware vSphere or Microsoft Hyper-V – will have an export function that creates the OVF packages based on the existing Virtual Machine images.

If your organization does not currently do any Virtualization, then you have two options to migrate to the VM based Cloud environment.

- Perform Physical to Virtual (P2V) VM images of each server
- Re-build each server in the new environment.

P2V based migrations rely on software tools that you run on your existing servers that produce a Virtual Machine image. P2V conversion tools are easily accessible on the web and are often free. The VM image produced will include the Operating System (OS), any installed applications, and all corresponding configuration changes. It is an exact image of your hardware server, but built in a VM file structure. P2V conversion is typically used when the server application stack is complex and has many customized configurations. It is often easier to do the P2V conversion rather than try and re-build the application stack and remember all the configuration tweaks that have been made over time.

Alternatively, rebuilding your application stack by installing new OS, patches and security configurations enables the opportunity to create a new stable environment for your new VMs. This is your opportunity to start from scratch and build an optimized environment. Using VM templating techniques the replication of this stable VM foundation can easily be used for all your applications within the Cloud platform. The appropriate approach for your specific environment will be identified during a “Virtualization Assessment” study.

As you can see there are a number of key issues to consider throughout the migration of your server footprint to the Cloud. Fortunately, once this migration is complete, the level of ongoing internal IT support will be reduced. Many IT intensive roles such as monitoring, backup, patching etc. are bundled into the Cloud offering. In-house IT resources are freed up to focus on proactive projects, introducing new business applications, or building new IT capabilities such as Test and Development environments. Financial resources are also freed up, as you move away from large up-front capital investments to a pay-per-use opex investment model.

## COMPLEX ►

# Build a hybrid Cloud and on-premise environment.

Hybrid IT models – A mix of on-premise & Cloud based infrastructure – are typically created for two purposes:

- To build flexible R&D or Test and Development sandboxes in the Cloud.
- To build replicated or mirrored Disaster Recovery (DR) sites in the Cloud.

The former scenario can be turned up easily and enables rapid turn-up of IT resources for development project, testing, or temporary campaigns. Contrary to the “Complex” nomenclature, this option is very easy to deploy as no synchronization is needed between the on-premise and Cloud infrastructures.

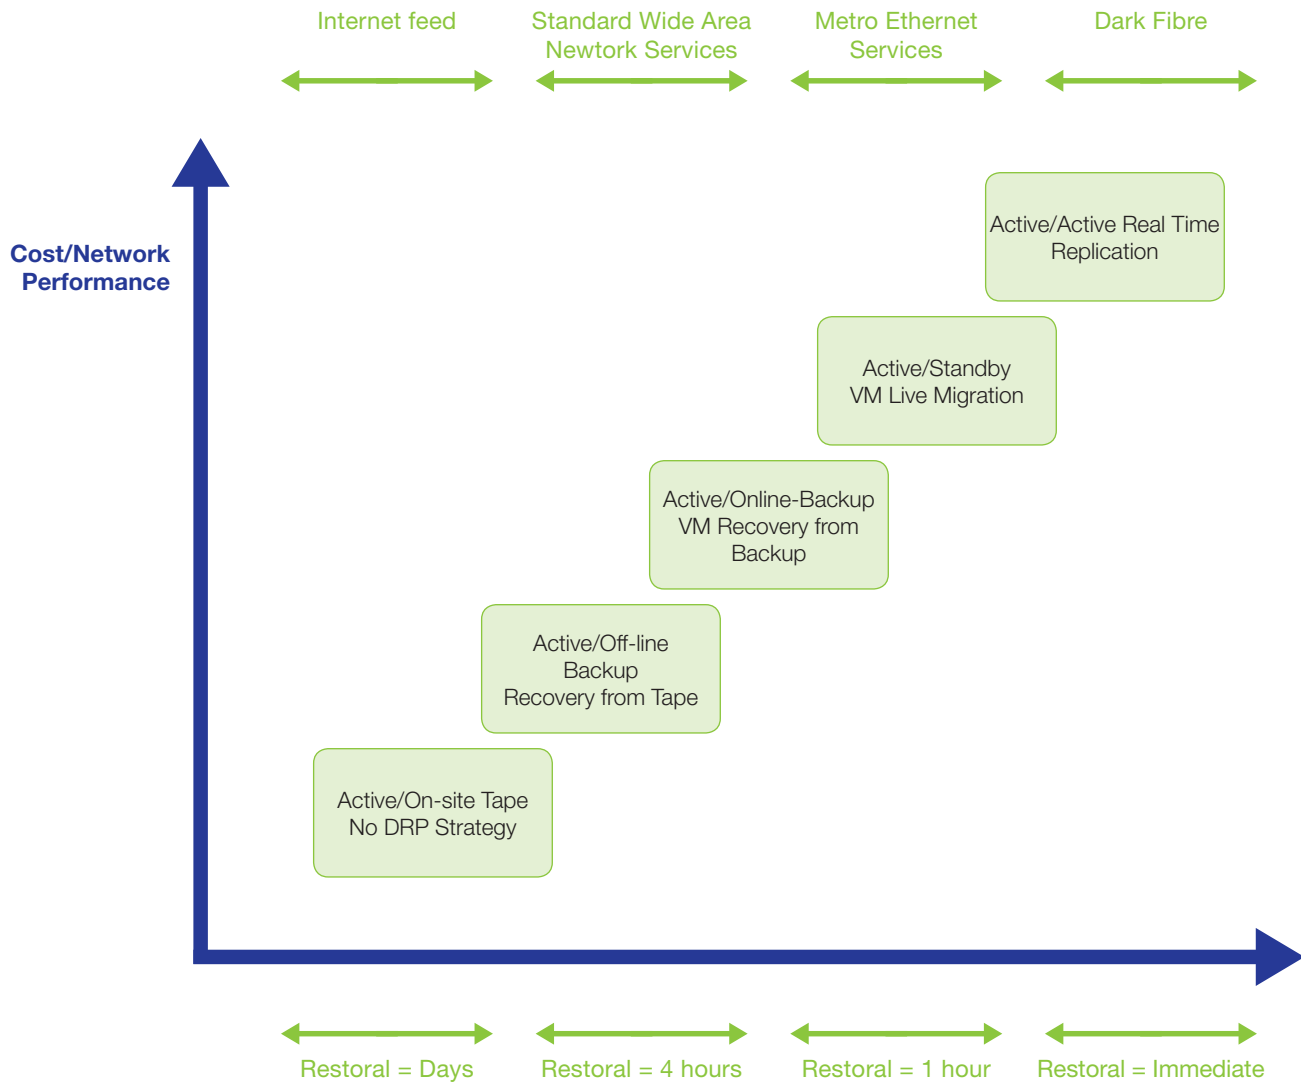
The “Complex” categorization applies to the second scenario where Hybrid IT is driven by Disaster Recovery or Business Continuity planning. The Cloud server farm is built as a standby IT environment for unexpected situations including fire or power outages, where the primary data centre becomes in-active.

This model is extremely complex – and often expensive – because of the synchronization required between the on-premise IT resources and those in the Cloud. High bandwidth levels (> 1 Gbps) and low latency levels (< 15 ms roundtrip) are needed for any Active/Active or Live VM Migration based DR strategies. The Gartner Group Study, “Assessing the Top Use Cases and Provider Types for Cloud-Based Disaster Recovery,” found the median quote for a 50 VM Cloud based DR solution was \$11,724 per month. Pricing for this configuration could be as high as \$50,000 monthly.

The scope of responsibility for in-house IT rises dramatically as they need to manage both environments; the production environment and the Cloud-based recovery environment. They must track inter-site mirroring or backups, keep VM images in sync, and periodically test their recovery processes. Due to the complexity, it is highly recommended to use DRP automation tools such as VMware Site Recovery Manager (SRM) that script the DRP workflows and enable non-impact failover testing. These tools, and the periodic DR testing they facilitate, will give your IT team confidence that their DR Failover will work as-planned when an actual emergency arises.

The following graph shows different Cloud based DR strategies and the corresponding Time-to-Restore targets. As restoral times are reduced, the cost for the DR solution increases significantly. This increase is due to the associated cost of Data Centre Interconnect (DCI) bandwidth and network performance. Any real-time or live migration scenarios require extremely low latencies, forcing data centre sites to be within a 100 Km range. These inter-data centre network connections can be very expensive. Most Cloud providers have meet-me facilities through which you can set up high performance DCI connections (eg: WAN, Metro Ethernet, or Dark Fibre) connections to the Cloud IT infrastructure.

# Cloud based Disaster Recovery Options



If the complexity and cost of building a real-time Hot or Replicated site is overwhelming, it is suggested to get your "toes wet" on the disaster recovery front using simpler Cold or Warm Site topologies that leverage Cloud based backup solutions.

## Conclusion

The majority of organizations can benefit from migrating some or all of their IT infrastructure to the Cloud. The business continuity benefits alone mean that Cloud based IT strategies deserve a second look. You should consider Cloud-based scenarios whenever business case modelling for upcoming technology refresh capital expenditures. In many cases the flexible pay-per-use opex based financing allows for a technologically superior architecture that is un-attainable if purchased outright.

From a deployment perspective, it is recommended to take a step-by-step approach and start with less critical workloads (eg: test and development servers) or self-contained Software-as-a-Service domains (eg: Email or CRM). In most cases the Medium complexity option of building high availability (HA) systems in a Cloud data centre provide the best uptime value per budget dollar spent. The fault-resiliency and Service Level Agreements (SLA) guaranteed by Cloud providers is a significant improvement over business continuity levels achievable within most on-premise server closets. Ninety percent of disaster scenarios are mitigated by this topology.

Hybrid IT (Complex) Cloud based topologies can be exponentially more expensive and resource intensive to maintain if real-time synchronization is required between the two environments. Further, depending on the business criticality of your key systems, zero lost data and immediate time to restores may be necessary. In these cases, the Cloud based hot-site is often cheaper than building two on-premise data centre environments.

In all scenarios, the flexible nature of Cloud services paired with the pay-as-you-go subscription models will mitigate risk for your organization's cloud adoption projects.

## About the Author

Bryan Ting is a Customer Systems Engineering Specialist at TELUS and focusses on all areas of Information Technology: Networking, Unified Communications, Security and Cloud & Hosted services. With over 15 years in the IT services industry, Bryan has worked with clients ranging from large Enterprise and Government to Small and Medium Business.

Bryan's passion is technology roadmapping with clients and performing live demonstrations of Information Communication Technologies (ICT) in the various TELUS Innovation Centres or Innovation Centre on Wheels demonstration centres

