

# Scalability and Economics of Citrix XenApp and Citrix XenDesktop 7.6 on Amazon Web Services

December 2014

## Authors

The following authors contributed to the creation of this deliverable.

### Citrix

Peter Bats  
Largo Augusto 8  
Milan, Italy  
Peter.Bats@citrix.com

## Revision History

Revision	Change Description	Updated By	Date
0.9	Initial Draft	Peter Bats	24 November 2014
0.99	Final Draft	Peter Bats	9 December 2014

# Table of Contents

- Authors ..... 2
- Revision History ..... 2
- Introduction ..... 5
  - Amazon Web Services ..... 6
  - What is hybrid cloud provisioning? ..... 6
    - Benefits of hybrid cloud provisioning ..... 7
  - FlexCast Management Architecture (FMA) models ..... 7
  - XenApp ..... 7
  - XenDesktop ..... 8
  - Private cloud versus public cloud ..... 8
  - Key use cases ..... 8
- Results summary ..... 10
  - XenDesktop AWS sample environments ..... 14
    - Sample site for 1,000 XenDesktop and XenApp users ..... 14
    - Sample site for 10,000 XenDesktop and XenApp users ..... 16
- Scalability testing in a multiuser environment ..... 19
- Test results ..... 19
  - Test setup and configurations ..... 19
    - EC2 instance test configuration for single instance scalability ..... 19
    - XenApp 7.6 configuration for 1,000 user scalability ..... 21
  - Testing methodology ..... 22
    - Load Generation ..... 22
    - User workload simulation—Login VSI from Login Consultants ..... 22
    - Testing procedure ..... 23
    - Success criteria ..... 24
- Test results: single instances ..... 25
  - Single-instance validation ..... 25
  - Compute Optimized C3 instances ..... 40
- Test results: Multi-instance validation—1,000 Users ..... 45

Scalability considerations and guidelines.....	48
AWS Instance configuration .....	48
AWS Storage best practices .....	48
How to change the UserProfile directory to point to drive E:.....	49
Appendix A: Performance counters .....	50
Appendix B: Cost per user for different workloads overview.....	52
Appendix C: Online resources .....	58
Amazon reference documents .....	58
Citrix reference documents .....	58

## Introduction

The emergence of cloud computing has changed every aspect of enterprise IT service delivery. IT organizations around the world have already started to embark on the journey to take their services “to the cloud.” This approach enables IT to leverage virtualized datacenter resources to gain automated, flexible, self-service clouds that ensure the best security, performance and reliability whether they are running in the enterprise datacenter, an external cloud, or a hybrid of both.

The virtualization movement and subsequent cloud-era have dramatically changed the way IT protects and delivers applications, desktops, and data to an ever-increasing mobile workforce. Increased adoption of virtualization technologies and cloud-based solution architectures brings about a myriad of benefits, but it also introduces additional layers of complexity for each new variation in the provisioning and delivery process.

Advancements in virtualization have added more layers as both physical servers, virtual servers and associated storage now need to be configured, managed and monitored independently across private, public or hybrid clouds. To reap the larger benefits, the once simple service delivery workflow has been expanded to account for the additional layers (including physical or virtual) and locations (including private, public or hybrid cloud) that are now integrated within the solution architecture.

The attractiveness of a Citrix XenApp® and XenDesktop® cloud-based application and desktop offering is driven by the ever-evolving demands of workforce growth, flexibility and geographical disbursement. The ability to reallocate upfront costs of a large physical hardware infrastructure purchase into an operating expense model distributed over an extended timeframe with on-demand instant access to resources is an extremely attractive proposition. The instant availability of cloud-based resources dramatically reduces the implementation time previously allocated for procuring, racking and cabling physical servers. The scalability and flexibility of addressing a dynamic shift in workforce—whether through a merger, acquisition or simple seasonal demand—provide even more business justification for a cloud-oriented deployment.

This paper explores the scalability and economics of hybrid cloud provisioning using XenApp and XenDesktop, utilizing their native cloud orchestration and provisioning capabilities on Amazon Web Services (AWS). The scalability and economics of such a hybrid cloud provisioning, through XenApp and XenDesktop utilizing Amazon Web Services, allow for delivery of virtual applications and desktops by enterprises or service providers with a price point as low as \$15 a month for VDI deployments or \$12 a month for XenApp deployments (less than a penny per user per hour, all Citrix and Microsoft licensing included).

## Amazon Web Services

Amazon Web Services (AWS) offers a broad set of global compute, storage, database, analytics, application, and deployment services that help organizations move faster, lower IT costs, and scale applications. These services are trusted by the largest enterprises and the hottest start-ups to power a wide variety of workloads including web and mobile applications, data processing and warehousing, storage, archive, and many others. These services also allow Citrix in collaboration with AWS to deliver its innovative desktop virtualization, networking, and enterprise mobility solutions on AWS.

AWS is cost effective, dependable, flexible and comprehensive. With AWS, you pay only for what you use, with no up-front expenses or long-term commitments. The Amazon cloud is scalable, with massive compute capacity and storage. It is reliable, redundant and secure.

Amazon Virtual Private Cloud (VPC) lets you provision a private, isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define. With Amazon VPC, you can define a virtual network topology that closely resembles a traditional network that you might operate in your own datacenter. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.

You can easily customize the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your webservers or NetScaler VPX™ virtual appliances that have access to the Internet and place your backend systems, such as databases, XenDesktop Delivery Controllers, or XenApp application servers, in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network-access-control lists, to help manage access to Amazon Elastic Compute Cloud (Amazon EC2) instances in each subnet.

## What is hybrid cloud provisioning?

Today many enterprises are increasingly turning to hybrid clouds, allowing them to combine the benefits of building private and public clouds as well as leverage existing IT infrastructure to cut costs, maximize value and modernize the way IT services are delivered. Hybrid clouds are formed when private and public clouds are closely integrated, delivering IT services with the security, control and agility users demand.

Citrix XenDesktop and XenApp have been redesigned as cloud-ready solutions for delivering both applications and desktops on any type of cloud infrastructure by integrating with any virtual infrastructure technology, storage infrastructure and complex network topologies to deliver a single, unified platform. XenDesktop and XenApp now enable enterprise IT to build a common service delivery architecture for all Windows apps and desktops leveraging common policies and tools that simplify deployment and management.

By delivering Windows apps and desktops as a cloud-like service, XenDesktop and XenApp can handle multiple versions and instances of both Windows Server and desktop operating systems from a single platform. XenDesktop and XenApp are built to leverage any virtual infrastructure or cloud management platform. Virtual apps and desktops can be deployed on popular cloud platforms including Apache CloudStack-based Citrix CloudPlatform™ or Amazon Web Services (AWS), making it easier than ever to dynamically expand the infrastructure footprint.

## Benefits of hybrid cloud provisioning

Through cloud integration XenDesktop and XenApp simplify the hardware and storage sizing and planning process by allowing IT administrators to deliver both applications and desktops from a single instance to:

- Simplify smaller deployments or span deployments across a variety of private, public, and hybrid clouds
- Quickly expand the infrastructure footprint and put less restrictions on upfront planning and sizing
- Easily scale down oversized environments and scale up undersized environments to reduce costs
- Enable administrators to have the infrastructure they need to deliver a high-performance user experience

## FlexCast Management Architecture (FMA) models

Employees across the enterprise have varying performance, personalization and mobility requirements. Some require offline mobility; others need simplicity and standardization; power users need a high-performance, fully personalized desktop; and almost every employee needs simple, secure access to a Windows app from their tablet, smartphone or laptop. Together XenDesktop and XenApp meet all these requirements with FlexCast<sup>®</sup> Management Architecture designed for increased scalability and flexibility. With XenDesktop and XenApp, IT can deliver every type of virtual desktop or app, hosted or local, optimized to meet the performance, security and mobility requirements of each individual user while optimizing ongoing management and deployment costs.

## XenApp

XenApp provides the following capabilities:

- **Secure, mobile access to Windows apps.** XenApp takes enterprise Windows app mobile by centralizing mission-critical business apps in the datacenter or the AWS cloud and delivering secure remote access on any device, anywhere. Taking Windows apps mobile and to the cloud has never been easier: XenApp can dynamically recognize a mobile device and automatically transform the application display for native mobile device features including touch-friendly menus, finger-swipe scrolling and pop-up controls. Even highly complex 3D graphical apps from the manufacturing, design, engineering, and construction industries can be accessed on tablets and smartphones with HDX™ 3D Pro technology and Amazon EC2 G2 GPU instances. XenApp with HDX 3D Pro technology supports hardware-based GPU using the Amazon EC2 G2 GPU instances, allowing for sharing of OpenGL-based 3D professional graphics apps for smooth graphics performance and breakthrough deep compression technologies that maximize performance over low-bandwidth, high-latency networks.
- **XenApp published desktops (shared, server-based desktop).** Based on Remote Desktop Session Host (RDSH) technology, XenApp enables multiple user sessions to connect to a single server with access to an isolated instance of a Windows server desktop for the most cost-efficient, high-performance virtual desktop solution designed to meet the needs of the mainstream workforce.
- **Access to key applications.** Almost every organization has one or two applications that run the business. Access to these applications is critical if the organization wants to maintain its customers. Applications form the core functionality of every endpoint device. In order to create content, lookup information or get a job done, users need access to applications. Finding a way to deliver the right application to the right user is the goal of XenApp.

## XenDesktop

XenDesktop provides the following features:

- **Pooled Virtual Desktop Infrastructure (VDI).** Through XenDesktop central image management technology, administrators can develop and manage a single desktop OS instance and seamlessly, on-demand provision that one instance out to thousands of users. This technology dramatically simplifies desktop patching and management while allowing employees to access their virtual desktop from a variety of devices and locations.
- **Personalized VDI.** Unlike a Pooled VDI deployment where user changes and customizations are prohibited or discarded between sessions, Personal vDisk technology enables user personalization and customizations to persist between desktop sessions. It provides users with the customized and personalized desktop experience they demand combined with the storage efficiency, centralization and management benefits of Pooled VDI.
- **Offline VDI (XenClient®)—Desktop virtualization to go with XenClient.** While the goal of leveraging virtualization is to centrally manage and host desktops and apps in the datacenter, there are cases where employees must be able to view and modify documents or data when disconnected from the network and offline. In these cases, XenClient permits administrators to stream and synchronize an entire managed desktop OS down to a local computer so that it can be taken offline as a complete encrypted file system with powerful policy enforcement.
- **Mobile access to pre-existing virtual or physical desktops with Remote PC Access.** Citrix delivers the most flexible desktop virtualization solution in the marketplace by supporting remote access to the pre-existing PCs in the workplace, as well as virtual desktops in the datacenter or cloud using the same broker, gateway appliance, and universal client components.

## Private cloud versus public cloud

The integrated platform for provisioning hybrid clouds using XenDesktop and XenApp enables IT administrators to deliver a complete range of apps and desktops, whether in the private cloud or in the public cloud, while consolidating management, monitoring and maintenance tasks.

- The private cloud is the collection of on-premise infrastructure, desktops, applications and data delivered on demand by enterprise IT. Private clouds can also be hosted off-premise. In this case, a service provider offers a portion of its public infrastructure for exclusive use by a single customer, also known as a tenant.
- The public cloud is the collection of off-premise, multitenant infrastructure, storage and computing resources, as well as Software as a Service (SaaS) applications and data, which are delivered on demand by external cloud service providers. Public clouds allow multiple customers, or tenants, to share the underlying resources with each paying only for the resources it consumes.

Users demand a seamless, reliable, high-performance experience from the desktop to the datacenter to the cloud. Citrix meets these demands by powering mobile work styles through solutions that address people, data and apps while powering cloud services through solutions for unifying, bridging and building cloud environments.

## Key use cases

Enterprise IT organizations that follow a hybrid cloud strategy get the best of both worlds by selecting which applications and usage scenarios fit best in their private cloud and which fit best in a public cloud, enabling them to flex, grow and transform to meet the demands of the modern workplace.

On-premise private clouds rely on capital investments in datacenters and infrastructure. Once these investments have been made, the incremental cost of running an application in a private cloud is often substantially less than running an equivalent workload on a public cloud. From an ROI perspective, this makes it important to achieve and maintain a high level of resource utilization in private clouds before moving applications to a public cloud.



Private clouds are also commonly used when direct control is needed. While private clouds don't necessarily have greater reliability or security than public clouds, governance requirements sometimes mandate that sensitive data, as well as applications that access that data, remain on premise. Some applications may also have dependencies on shared IT services or shared data that must remain on premise.

Public clouds are well suited for a number of other scenarios, particularly when the private cloud is operating at capacity. The underlying concept is based on owning just enough infrastructure so that it can be kept busy consistently. Then, rather than owning more infrastructure that would often sit idle, additional demands can be met through a public cloud. In this scenario, applications that only run periodically are great candidates to run in a public cloud. Test, demo and training environments aren't used continuously, so they also fit well within public clouds.

Public clouds provide seemingly unlimited scale, making them appropriate for situations with unpredictable demand. For example, product launches and promotional events can generate surprisingly high website traffic. These can be some of the most important times for business success, and IT must be ready to serve all prospective customers. Similarly, seasonal demands—from holiday shopping to tax season—may be best met by taking advantage of public cloud resources.

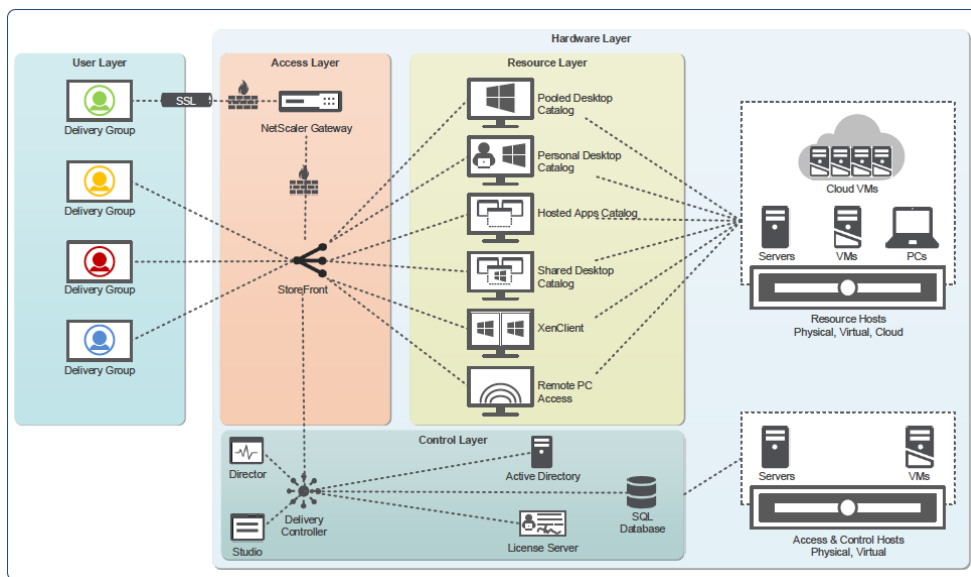


Figure 1. Architectural overview of XenApp and XenDesktop deployment.

Delivering these desktops and applications from XenDesktop and XenApp, as shown in Figure 1, presents unique challenges to administrators who are accustomed to delivering traditional cloud-hosted applications or services. The ability to determine how many users can be hosted on a single Amazon Elastic Compute Cloud (Amazon EC2) instance is critical to achieve maximum operation efficiency.

This document will guide you through the selection and sizing of the different Amazon EC2 instance types, as well as the economic impact of using these instances, allowing you to further reduce cost while securely delivering applications to anyone, anywhere, anytime, on any device.

# Results summary

In order to determine user scalability of XenApp servers, the Login VSI performance and scalability testing tool from Login Consultants was used to simulate real-world user interactions with two different sets of configurations: Windows Server 2012R2 with Microsoft Office Professional 2013, and Windows Server 2008R2 with Microsoft Office Professional 2010 applications. Login VSI allows for complete test automation including session launch, user action simulation, and the monitoring of session performance characteristics.

The bar graphs below (Figure 2 through Figure 5) shows the maximum number of simulated XenApp user sessions for the Task Worker workload and the Knowledge Worker workload that each EC2 instance configuration can support before encountering a user-experience bottleneck. In the same graphs you'll also find the cost per user for these EC2 instance types. Additional details can be found in the accompanying tables (Table 1 and Table 2). Note especially the t2.small and t2.medium instance types with respectively a \$0.006 and a \$0.007 per user per hour compute cost. Please see also the detailed results section for more information on the bottlenecks as well as a description of these workloads.

Figure 2 and Figure 3 show results for Windows Server 2012R2 and Office 2013.




Figure 2. Knowledge Worker Workload users per EC2 instance type: XenApp 7.6, Windows Server 2012R2, and Office 2013.




Figure 3. Task Worker Workload users per EC2 instance type: XenApp 7.6 – Windows Server 2012R2 – Office 2013.

Figure 4 and Figure 5 show results for Windows 2008R2 and Office 2010.




Figure 4. Knowledge Worker Workload users per EC2 instance type: XenApp 7.6 – Windows 2008R2 – Office 2010.




Figure 5. Task Worker Workload users per EC2 instance type: XenApp 7.6 – Windows 2008R2 – Office 2010.

Table 1. Task Worker CCUs per current instance type overview.

Task Worker Workload - XenApp 7.6 - Windows Server 2012R2 - Office 2013							
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	Cost per User/h	
<b>General Purpose</b>							
t2.micro	Var	1	1	\$ 0.018	1	\$	0.018
t2.small	Var	2	1	\$ 0.036	6	\$	0.006
t2.medium	Var	4	2	\$ 0.072	10	\$	0.007
m3.medium	3	3.75	1	\$ 0.133	4	\$	0.033
m3.large	6.5	7.5	2	\$ 0.266	9	\$	0.030
m3.xlarge	13	15	4	\$ 0.532	19	\$	0.028
m3.2xlarge	26	30	8	\$ 1.064	42	\$	0.025
<b>Compute Optimized</b>							
c3.large	7	3.75	2	\$ 0.188	10	\$	0.0188
c3.xlarge	14	7.5	4	\$ 0.376	22	\$	0.0171
c3.2xlarge	28	15	8	\$ 0.752	42	\$	0.0179
c3.4xlarge	55	30	16	\$ 1.504	83	\$	0.0181
c3.8xlarge	108	60	32	\$ 3.008	164	\$	0.0226
<b>Memory Optimized</b>							
r3.large	6.5	15	2	\$ 0.300	9	\$	0.033
r3.xlarge	13	30.5	4	\$ 0.600	18	\$	0.033
r3.2xlarge	26	61	8	\$ 1.080	35	\$	0.031
r3.4xlarge	52	122	16	\$ 1.944	69	\$	0.028
r3.8xlarge	104	244	32	\$ 3.500	135	\$	0.026
<b>Storage Optimized</b>							
i2.xlarge	14	30.5	4	\$ 0.973	18	\$	0.054
i2.2xlarge	27	61	8	\$ 1.946	35	\$	0.056
i2.4xlarge	53	122	16	\$ 3.891	69	\$	0.056
i2.8xlarge	104	244	32	\$ 7.782	135	\$	0.058

## XenDesktop AWS sample environments

For reference, two types of XenDesktop AWS sample environments are provided for an estimate of the costs of running a XenDesktop and XenApp 7.6 platform on AWS. These sample reference designs are for a 1,000 Concurrent Users (CCU) and 10,000 CCU XenDesktop and XenApp site, respectively.

Both designs use a core infrastructure that is active and powered-on 24 hours a day (Delivery Controller, SQL Database, NetScaler Gateway™), and a dynamic, elastic part. The multisession XenApp servers are active for 10 hours a day, covering normal 8-hour work days as well as providing ample margin for workshifting. The single-session personal desktop VDI instances are active and powered-on for 8 hours a day. The XenDesktop Delivery Controller will automatically power up and down the VDI instances on-demand during the day. Note also that the core infrastructure is deployed in a distributed redundant fashion across multiple Availability Zones, to reduce the impact of any outages at the infrastructure layer.

The design makes use of a one-year term of AWS Reserved Instances for the core infrastructure and on-demand instances for the dynamic, elastic part. With Amazon EC2 Reserved Instances you pay a low, one-time fee and in turn receive a significant discount on the hourly charge for that instance.

See <http://aws.amazon.com/ec2/reserved-instances/> on the Amazon Web Services site for more details. Because the dynamic instances are powered on for a relatively low number of hours per month, they are not ideal candidates for the usage of reserved instances as well.

The XenApp or XenDesktop worker instances in this design have been configured with Amazon Elastic Block Store (EBS) storage for their boot device and with local instance storage to hold volatile data such as pagefiles and user profiles. EBS storage provides block-level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persist independently from the life of an instance. Amazon EBS provides high-availability and highly reliable storage volumes that are placed and automatically replicated in a specific Amazon Availability Zone.

Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low-latency network connectivity to other Availability Zones in the same Region. By launching instances in separate Availability Zones, you can protect your applications from failure of a single location. Regions consist of one or more Availability Zones, are geographically dispersed, and will be in separate geographic areas or countries. Citrix Studio, the management console of XenDesktop and XenApp, is tightly integrated with the EC2 infrastructure and allows for direct deployment in the required Availability Zones.

With Amazon EBS, volume storage is charged by the amount of storage allocated and is priced at a rate of \$0.05 for magnetic disks and \$0.1 for SSD disks per allocated GB per month. Amazon EBS, however, also charges \$0.05 per 1 million I/O requests to a magnetic disk volume. No charges are applied for I/O requests to SSD disk volumes.

Given that each XenApp session generates on average 4 IOPS per user (see the Citrix blog post “What’s the optimal XenApp 6.5 VM configuration?” at <http://blogs.citrix.com/2013/01/07/whats-the-optimal-xenapp-6-5-vm-configuration/>), it is advised that you use local-instance storage instead of EBS for any volatile data. Local-instance storage is an inexpensive way to provide high-speed storage and only persists during the life of the instance. This strategy lowers the amount of IOPS per user to as low as 1.5 IOPS per user, as can be seen in the test results (see Figure 45. C Drive IOPS versus disk queue length for c3.4xlarge Task Worker workload., which shows a maximum amount of disk transfer/sec (IOPS) to the EBS system drive of 134 IOPS for 84 active sessions.

### Sample site for 1,000 XenDesktop and XenApp users

Figure 6 shows an example site on AWS that supports 1,000 concurrent hosted apps and desktops using XenDesktop and XenApp.

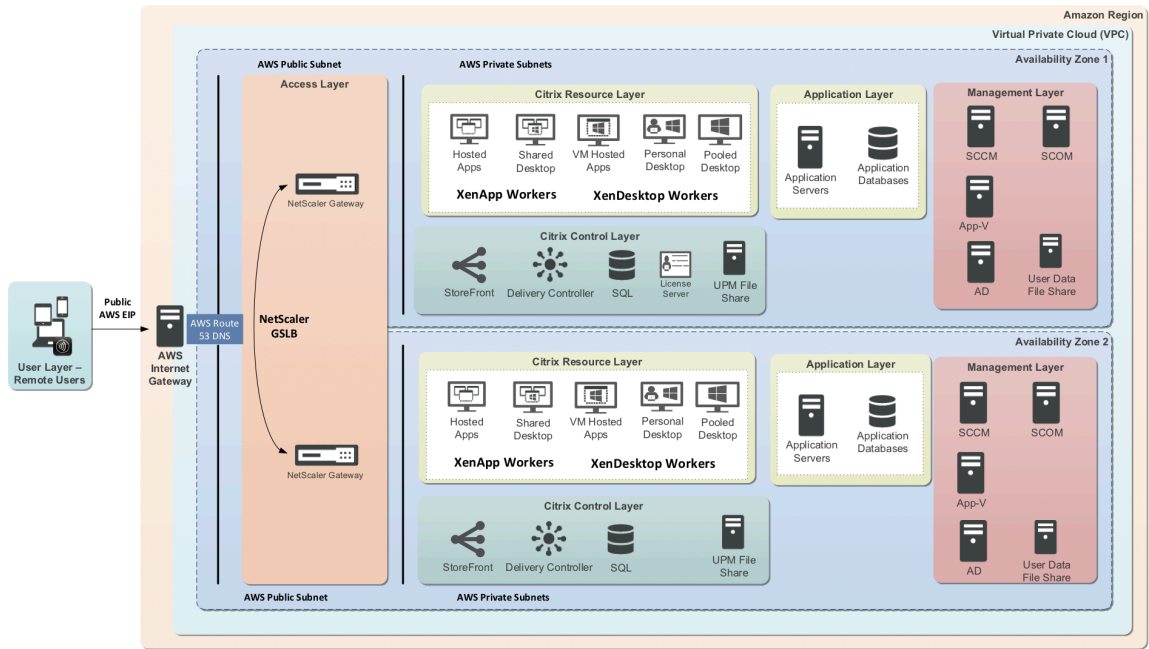


Figure 6. XenApp and XenDesktop On-Cloud Site and AWS—supports 1,000 hosted apps and desktops.

Figure 7 shows the corresponding costing information for this 1,000 desktop example site on AWS.

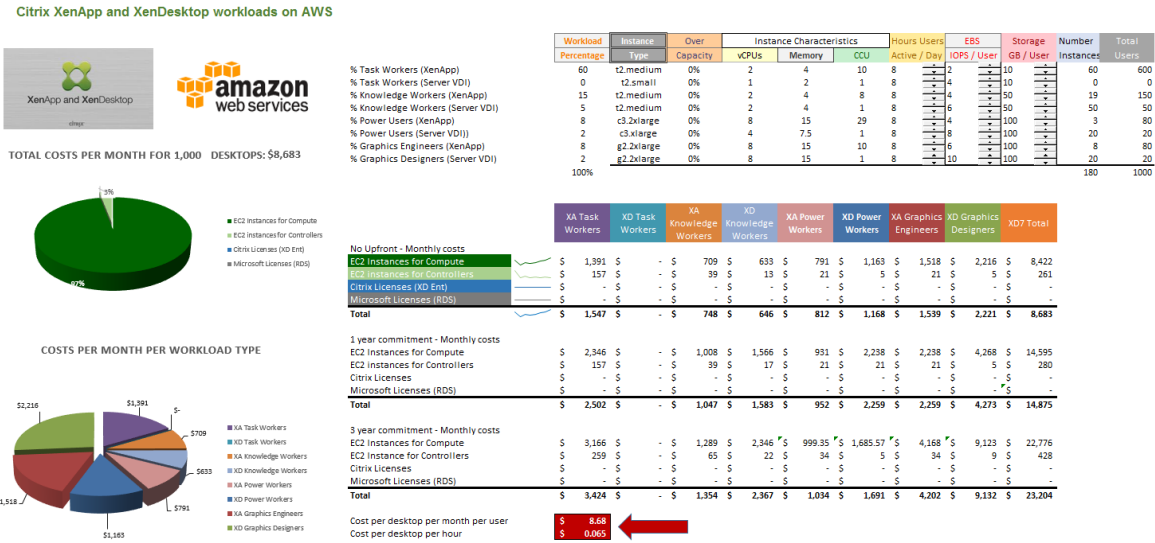


Figure 7. XenDesktop and XenApp on AWS—1,000 desktops costing information example.

This sample design for 1,000 users composed of a mix of various workload types makes use of two NetScaler Gateway™ instances (m3.xlarge) to provide access, one per AWS Availability Zone. It uses one XenDesktop Delivery Controller/StoreFront (c3.large), one per Amazon Availability Zone, as well as one SQL Server instance (r3.large) per Availability Zone to host the various XenDesktop

site databases. The XenDesktop License Server is hosted on one of the XenDesktop Delivery Controllers.

The example for 1,000 users assumes a mix of workloads and FlexCast models to best meet the demands of the various users. In this example, 600 users (60%) are considered Task Workers using the XenApp Remote Shared Desktop model; 200 users (20%) are considered Knowledge Workers, of which 2% require a dedicated personal VDI workstation. Another 100 users (10%) are Power Users of which 20 require a powerful 4 vCPU, 8GB dedicated personal VDI instance. And a final 100 users (8%) are graphical engineers and designers that require access to hardware GPU acceleration; of which 20 users (20%) are graphics designers that need a dedicated GPU.

For each workload type, an appropriate EC2 instance type has been selected.

Workload	Instance Type	Instance Characteristics			EBS IOPS / User	Storage GB / User	Number Instances	Total Users	Cost per user per hour
		vCPUs	Memory	CCU					
% Task Workers (XenApp)	t2.medium	2	4	10	2	10	60	600 \$	0.019
% Task Workers (Server VDI)	t2.small	1	2	1	4	10	0	0	
% Knowledge Workers (XenApp)	t2.medium	2	4	8	4	50	19	150 \$	0.037
% Knowledge Workers (Server VDI)	m3.large	2	7.5	1	6	50	50	50 \$	0.112
% Power Users (XenApp)	c3.2xlarge	8	15	29	4	100	3	80 \$	0.070
% Power Users (Server VDI)	c3.xlarge	4	7.5	1	8	100	20	20 \$	0.281
% Graphics Engineers (XenApp)	g2.2xlarge	8	15	10	6	100	8	80 \$	0.070
% Graphics Designers (Server VDI)	g2.xlarge	4	7.5	1	10	100	20	20 \$	0.281
100%							180	1000 \$	0.124

Figure 8. Spreadsheet for calculating estimated costs (1000 users).

The example of 1,000 users of mixed workloads and FlexCast models used above would have a cost per desktop per month of \$8.68, or \$0.065 per user per hour, excluding any Citrix XenDesktop or Microsoft RDS SAL licenses. Including the Citrix XenDesktop Enterprise and Microsoft RDS SAL licenses, using public US list prices would bring us to \$16.84 per desktop per month or \$0.126 per user per hour.

A spreadsheet (see Figure 8) has been made available as a tool to assist in calculating your own estimated costs for running your XenApp and XenDesktop infrastructure on Amazon Web Services. It also allows for adding the cost of the needed Microsoft Remote Desktop Services (RDS) Subscriber Access License (SAL) and Citrix XenDesktop licenses to be able to have a more complete picture, reflecting the actual cost and discount models in use for those licenses.

The spreadsheet can be downloaded at <https://s3.amazonaws.com/cf-XenApp/DaaS+on+Public+Cloud+Costing+IaaS+-+AWS+-+all.xlsx>.

### Sample site for 10,000 XenDesktop and XenApp users

Figure 9 shows an example site on AWS that supports 10,000 concurrent hosted apps and desktops using XenDesktop and XenApp.



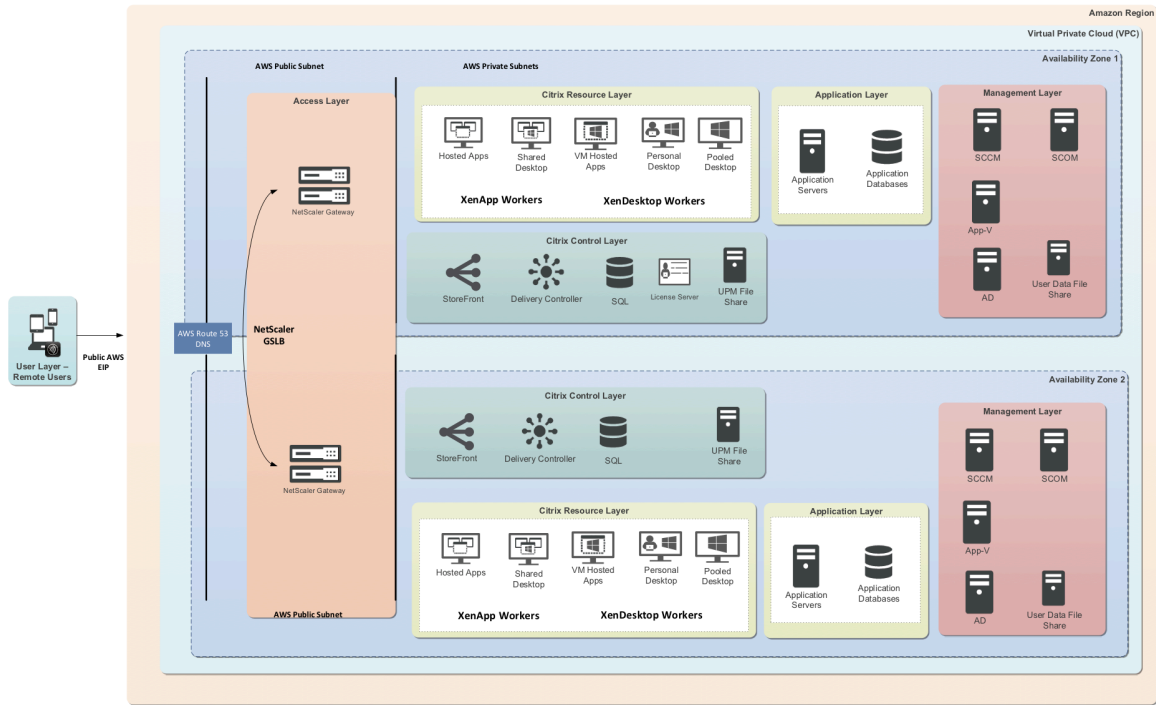


Figure 9. XenApp Farm on AWS for 10,000 hosted shared desktops.

Figure 10 shows the corresponding costing information for this 10,000 desktop example site on AWS.

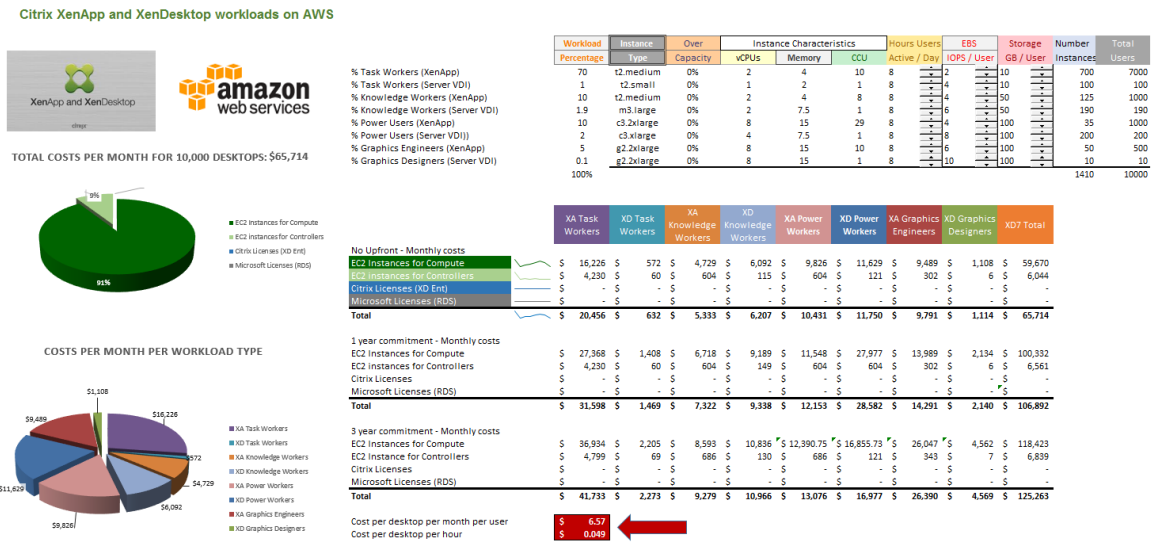


Figure 10. XenDesktop and XenApp on AWS 10,000 desktops costing information example.

This sample design makes use of the combination of AWS Route 53 DNS web service and Citrix NetScaler Global Server Load Balancing (GSLB) to provide access via its included NetScaler Gateway functionality, and uses two m3.2xlarge instances per Amazon EC2 Availability Zone. It uses four c3.xlarge instances, two per Amazon EC2 Availability Zone, to host the Delivery

Controller, StoreFront and License Server features. A SQL Database Server is set up in redundant mode using SQL Server 2012 AlwaysOn Availability Group or alternatively SQL Mirroring. This design is based on an r3.large instance per Amazon EC2 Availability Zone.

Below are the high availability solutions to consider for ensuring automatic failover on AWS:

- **SQL Mirroring** is the recommended solution. Mirroring the database makes sure that, should you lose the active database server, the automatic failover process happens in a matter of seconds, so that users are generally unaffected.
- **AlwaysOn Availability Groups** is an enterprise-level high-availability and disaster-recovery solution introduced in SQL Server 2012 to provide maximize availability for one or more user databases. AlwaysOn Availability Groups requires that the SQL Server instances reside on Windows Server Failover Clustering (WSFC) nodes. For more information, see “AlwaysOn Availability Groups (SQL Server)” at <http://msdn.microsoft.com/en-us/library/hh510230.aspx> and “Implementing Microsoft Windows Server Failover Clustering and SQL Server 2012 AlwaysOn Availability Groups in the AWS Cloud” at <http://aws.amazon.com/whitepapers/microsoft-wsfc-sql-alwayson/>.

This example for 10,000 users also assumes a mix of workloads and FlexCast models to best meet the demands of the various users. In this case, 7,100 users are considered Task Workers, with the majority (7,000) using the XenApp Remote Shared Desktop model and 100 users (1%) that require a dedicated personal VDI desktop. An additional 1190 users are considered Knowledge Workers, of which 190 require a 2 vCPU, 7.5GB dedicated personal VDI workstation (m3.large). An additional 1200 users are Power Users, of which 200 require a powerful 4 vCPU, 8GB dedicated personal VDI instance (c3.xlarge). And another 500 users are graphical engineers and designers that require access to hardware GPU acceleration, of which 10 users are graphics designers that need a dedicated GPU.

The example of 10,000 users of mixed workloads and FlexCast models used above would have a cost per desktop per month of \$6.57, or \$0.049 per user per hour (less than half a penny), excluding any Citrix XenDesktop or Microsoft RDS SAL licenses. Including the Citrix XenDesktop Enterprise and Microsoft RDS SAL licenses, using public US list prices would bring us to \$14.92 per desktop per month or \$0.112 per user per hour (a little more than a penny per hour).

As indicated in the example with 1,000 users, a spreadsheet (see Figure 11) has been made available as a tool to assist in calculating your own estimated costs for running your XenApp and XenDesktop infrastructure on Amazon Web Services. This spreadsheet enables you to add the cost of the needed Microsoft Remote Desktop Services (RDS) Subscriber Access License (SAL) and Citrix XenDesktop licenses to be able to have a more complete picture, reflecting the actual cost and discount models in use for those licenses.

This spreadsheet can be downloaded at <https://s3.amazonaws.com/cf-XenApp/DaaS+on+Public+Cloud+Costing+IaaS+-+AWS+-+all.xlsx>.

For each workload type an appropriate EC2 instance type has been selected. Figure 11 shows the spreadsheet that was used to calculate the estimated costs for 10,000 XenApp and XenDesktop users.

Workload	Instance	Instance Characteristics				EBS	Storage	Number	Total	Cost per user
		Type	vCPUs	Memory	CCU					
% Task Workers (XenApp)	70	t2.medium	2	4	10	2	10	700	7000	\$ 0.022
% Task Workers (Server VDI)	1	t2.small	1	2	1	4	10	100	100	\$ 0.047
% Knowledge Workers (XenApp)	10	t2.medium	2	4	8	4	50	125	1000	\$ 0.040
% Knowledge Workers (Server VDI)	1.9	m3.large	2	7.5	1	6	50	190	190	\$ 0.211
% Power Users (XenApp)	10	c3.2xlarge	8	15	29	4	100	35	1000	\$ 0.040
% Power Users (Server VDI)	2	c3.xlarge	4	7.5	1	8	100	200	200	\$ 0.200
% Graphics Engineers (XenApp)	5	g2.2xlarge	8	15	10	6	100	50	500	\$ 0.080
% Graphics Designers (Server VDI)	0.1	g2.2xlarge	8	15	1	10	100	10	10	\$ 4.000
	100%							1410	10000	\$ 0.580

Figure 11. Spreadsheet for calculating estimated costs (10,000 users).

# Scalability testing in a multiuser environment

Several important factors need to be taken into consideration when determining the number of users an instance can support. Sizing a XenApp or XenDesktop instance depends on the following criteria:

- Specifications of the instance type, such as CPU, memory, disk and network capacity
- Applications, their requirements and how they interact with the system
- Degree of user activity

Before any tests are executed, instance resource limits and response time thresholds need to be set in order to determine the number of users a particular instance can support. For example:

- Eighty percent CPU usage
- Seventy percent memory usage, up to the point where paging begins to saturate the disk
- Acceptable user response times ( $\leq 4000\text{ms}$ )

These limits should be set to accommodate utilization peaks and future growth; optimal user density of an instance does not equate to its maximum capacity. If you size a server based solely on its maximum capacity, you can expect problems such as instance instability and poor user response times. Thus, in each performance test, consider the relationship between instance performance and the user experience.

For example, high CPU utilization can be the result of a low memory condition where the system is swapping applications from memory to disk, causing a large number of disk interrupts. Without collecting data for all of the machine's major subsystems, we would easily conclude that the bottleneck was the CPU, when it was a memory bottleneck all along.

Please reference the chart in Appendix A: Performance counters for performance counters for all of the major subsystems (CPU, memory, disk and network) that should be monitored in these types of scalability tests.

## Test results

The following section gives a detailed performance analysis for the tests used to determine the scalability impact of the configuration of the various EC2 instances for single session (VDI) and multisession (RDSH) usage. The Login VSI performance tool used to conduct the scalability tests provides a mechanism to automatically collect user experience information (VSIIndex\_avg). In order to determine the number of users a configuration can support, this user experience data is overlaid with the performance data of the subsystem that experienced a performance bottleneck. The point where user experience and the subsystem bottleneck intersect determines the number of users that the configuration can support. For convenience, a chart with all subsystem data is displayed for each test.

The guidelines presented in this section should be used as a reference and are only valid if your workloads and applications are similar to those executed in our tests. As always, we recommend executing your own scalability tests, utilizing your selected instances and applications, when making your final sizing decisions.

## Test setup and configurations

In this project, we tested the different available Amazon EC2 instances as single instances as well as part of a XenApp 7.6 farm hosting 1,000 users to illustrate the linear scalability.

### EC2 instance test configuration for single instance scalability

Figure 12 shows the test configuration for single instance scalability.

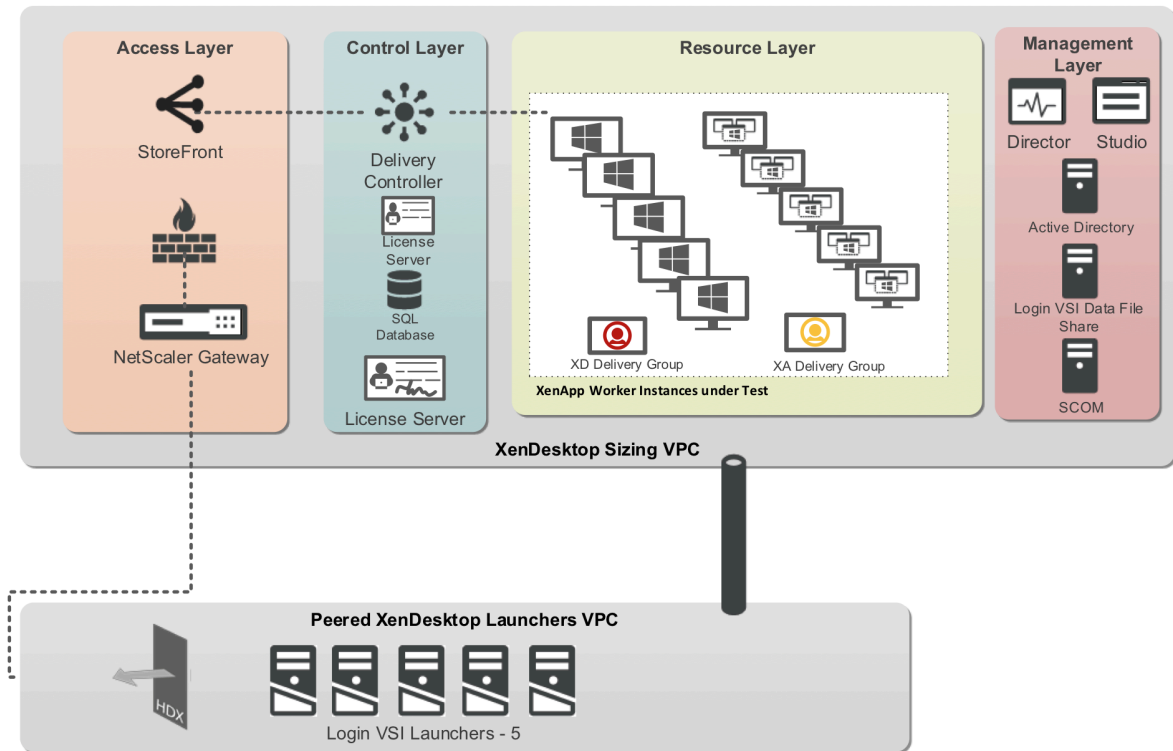


Figure 12. Configuration for the single-instance scalability test.

This test configuration contains the following AWS Instances:

- 1 x domain controller also functioning as DNS server
- 1 x Delivery Controller
- 1 x StoreFront
- 1 x VSI Share and controller
- 5 x t2.small VSI launchers
- EC2 Security Groups
- Private security group
- Public security group
- XenApp and XenDesktop VDI Instances configured with a single EBS boot volume and 1-2 local instance storage volumes, depending on the instance type. The local instance volumes are configured to hold volatile data such as the pagefile and user profiles.

This test configuration uses the following software components:

- Standard EC2 Windows Server 2012R2 or Windows Server 2008R2 AMIs 64-bit
- XenApp / XenDesktop 7.6 Delivery Controller and VDA
- Citrix StoreFront 2.6
- Citrix Receiver™ 4.1
- Microsoft Office 2013 or Microsoft Office 2010
- Login VSI 4.1

## XenApp 7.6 configuration for 1,000 user scalability

Figure 13 shows the configuration used for the 1,000 user scalability test.

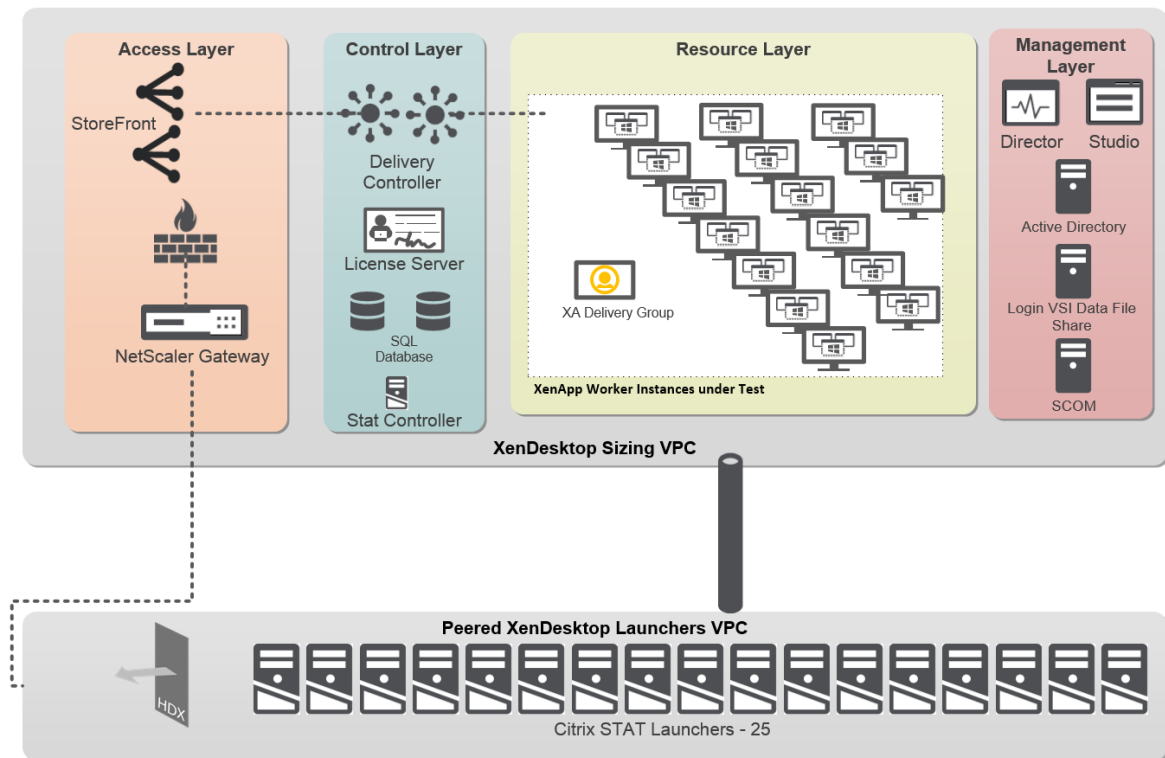


Figure 13. Configuration for the 1,000 user scalability test.

This test configuration contains the following AWS components:

- 2 x c3.xlarge domain controller, also functioning as DNS server
- 2 x c3.xlarge Delivery Controller
- 2 x c3.large StoreFront
- 25 x t2.small VSI launchers
- 16 x c3.4xlarge XenApp Workers
- EC2 Security Groups
- Private security group
- Public security group
- EBS Volumes for XenApp server instances
- 1 x 50GB EBS magnetic disk boot volume
- 1 x local instance ephemeral SSD disk for user profiles
- 1 x local instance ephemeral SSD disk for pagefile

This test configuration uses the following software components:

- Windows Server 2012R2 AMIs 64-bit
- XenApp 7.6 Delivery Controller and VDA
- Citrix StoreFront 2.6
- Citrix Receiver 4.1
- Microsoft Office 2013
- Login VSI 4.1

## Testing methodology

The objective of the test was to demonstrate XenDesktop 7.6 scaling capabilities and to understand the Amazon EC2 and EBS utilization.

Through conducting the same series of single server scalability tests against the different instance types, we were able to run consistent and comparable tests.

### Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the XenApp and XenDesktop 7.6 environment and executing a typical end-user workflow. To generate load within the environment, an auxiliary software application was required to generate the end-user connection to the XenDesktop environment, to provide unique user credentials, to initiate the workload, and to evaluate the end-user experience.

### User workload simulation—Login VSI from Login Consultants

One of the most critical factors of validating a XenApp/XenDesktop deployment is identifying a real-world user workload that is easy for customers to replicate and standardize across platforms in order to allow customers to realistically test the impact of a variety of worker tasks. To accurately represent a real-world user workload, a third-party tool from Login Consultants was used throughout our testing.

The tool has the benefit of taking measurements of the in-session response time, which provides an objective way to measure the expected user experience for individual desktop throughout large-scale testing, including login storms.

The Virtual Session Indexer, Login Consultants' Login VSI 4.1 methodology, the industry standard load testing tool for virtualized desktop environments, is completely platform and protocol independent and hence allows customers to easily replicate the testing results in their environment.

Login VSI 4.1, the version used in our scalability tests, comes with four preconfigured workloads but also allows users of the tool to create their own customized workloads (see Figure 14).

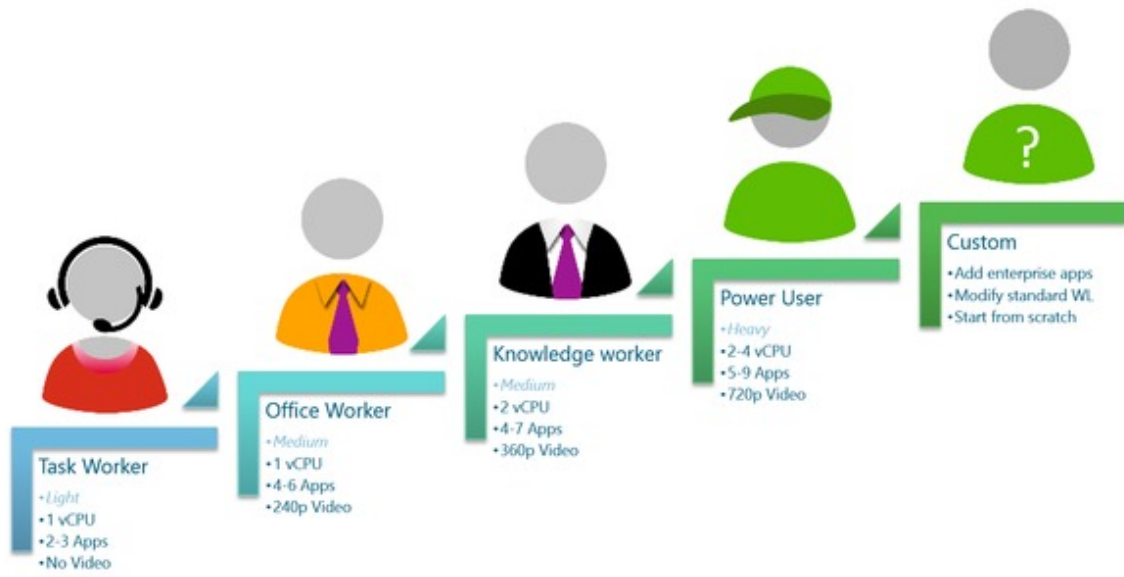


Figure 14. Login VSI 4.1 supports preconfigured and customized workloads.

Details on the four preconfigured Login VSI 4.1 workloads are included in Table 2.

Table 2. Login VSI 4.1 preconfigured workloads.

Workload	Profile Weight	vCPU	Apps	Video
Task Worker	Light	1	2-3	No
Office Worker	Medium	1	4-6	240p
Knowledge Worker	Medium	2	4-7	360p
Power User	Heavy	2-4	5-9	720p

Like real users, the scripted Login VSI session will open multiple applications at the same time. The Task Worker and Knowledge Worker workloads were used for this testing.

The Task Worker workload emulates a light knowledge worker using Excel, Outlook, Internet Explorer, Adobe Reader, Copy, Zip and printing actions.

The Knowledge Worker workload has been designed around a user utilizing at times two vCPUs. It is also the default workload in Login VSI. This workload emulates a medium knowledge worker using Office, Internet Explorer, PDF and Java/FreeMind.

You can obtain additional information on Login VSI from <http://www.loginvsi.com>.

## Testing procedure

The following protocol was used for each test cycle in this study to ensure consistent results.

### Pre-test setup for single-instance and complete farm testing

All instances under test were shut down utilizing Citrix Studio. All launchers for the test were shut down from the AWS Console. They were then restarted until the required number of launchers were running with the Login VSI Agent at a “waiting for test to start” state.



## Success criteria

Multiple metrics were captured during each test run, but the success criteria for considering a single test run as pass or fail was based on the key metric VSI Max. The VSI Max evaluates the user response time during increasing user load and assesses the successful start-to-finish execution of all the initiated virtual desktop sessions.

### Login VSI Max

Login VSI Max (VSI Max) represents the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users as indicated during the workload execution. The details for the exact calculation of VSI Max can be found here

<http://www.loginvsi.com/documentation/index.php?title=VSImax>.

If VSI Max is reached, that indicates the point at which the user experience has significantly degraded. The response time is generally an indicator of the host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.

### Determining VSI Max

The Login VSI analyzer will automatically identify the VSI Max. In the example below (Figure 15), the VSI Max is 98. The analyzer will automatically determine “stuck sessions” and correct the final VSI Max score.

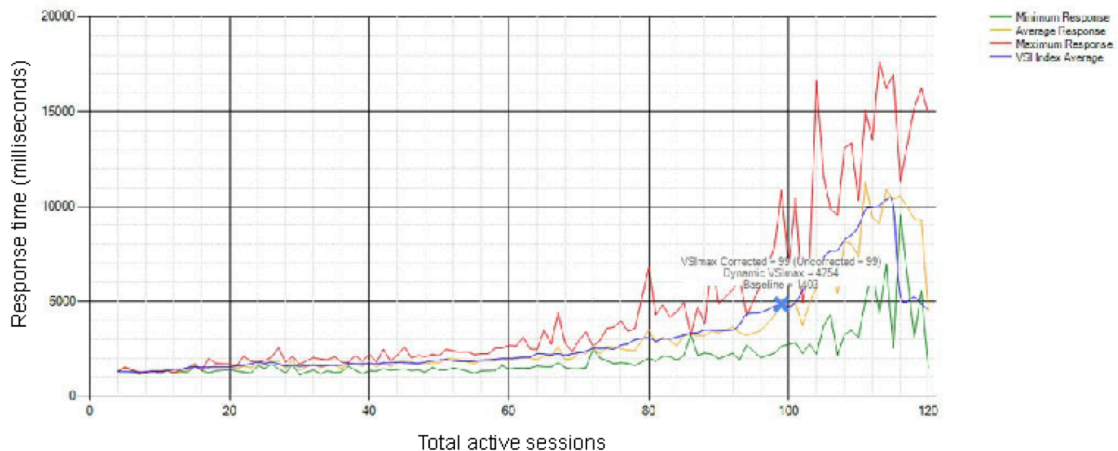


Figure 15. Sample Login VSI Analyzer graphic output.

In Figure 15, the following response times are shown:

- Red line: maximum response (worst response time of an individual measurement within a single session)
- Orange line: average response time within for each level of active sessions
- Blue line: the VSI Max average
- Green line: minimum response (best response time of an individual measurement within a single session)



The Citrix Director will be monitored throughout the steady state to insure that:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down.

## Test results: single instances

The purpose of this testing is to provide the data needed to validate Citrix XenApp and XenDesktop 7.6 Hosted Shared Desktop FlexCast model (see “Citrix Virtual Desktop Handbook” at [http://support.citrix.com/servlet/KbServlet/download/35949-102-713877/Citrix\\_Virtual\\_Desktop\\_Handbook\\_\(7x\).pdf](http://support.citrix.com/servlet/KbServlet/download/35949-102-713877/Citrix_Virtual_Desktop_Handbook_(7x).pdf) for a description on the various FlexCast models) on Amazon EC2 to virtualize Microsoft Windows Server 2008R2 or Microsoft Windows Server 2012R2 desktops, as well as the XenDesktop 7.6 Hosted VDI FlexCast model.

The information contained in this section provides data points that a customer may reference while designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined in this paper, and do not represent the full characterization of XenApp and XenDesktop on Amazon EC2.

For the Hosted VDI model, the test performed was to see that certain instance types were able to successfully complete Login VSI test suite for different defined workloads. A XenDesktop Delivery Group of five VDI machines was created for this purpose, and a successful test would mean that all five instances would have to complete the test for one single user. VSI Max is not used in this scenario as we’re not trying to establish how many Hosted VDI desktops the Amazon EC2 Cloud can support. See respectively the “W2K8R2 – O2K10” and “W2K12R2 – O2K13” worksheets in the accompanying spreadsheet which can be downloaded from S3 Amazon Web Services at <https://s3.amazonaws.com/cf-XenApp/DaaS+on+Public+Cloud+Costing+IaaS+-+AWS+-+all.xlsx>. The tests basically showed that all EC2 instance types were able to successfully complete the Login VSI test suite for the Task Worker and Knowledge Worker workloads. Only the smaller t2.micro, t1.micro and m1.small instance types were unable to successfully complete the Power Worker workloads.

This paper provides primary success criteria metrics for the most relevant to XenApp and XenDesktop EC2 instance types, an example of what is possible under the specific environment conditions outlined in this paper, and does not represent the full characterization of XenApp and XenDesktop while deployed on Amazon EC2.

### Single-instance validation

The following sections include the results for single-instance validation, including general purpose instances and compute optimized instances.

#### General purpose T2 instances

T2 instances are *Burstable Performance Instances* that provide a baseline level of CPU performance with the ability to burst above the baseline. The baseline performance and ability to burst are governed by CPU Credits. Each T2 instance receives CPU Credits continuously at a set rate depending on the instance size (see Table 3). T2 instances accrue CPU Credits when they are idle and use CPU credits when they are active. T2 instances are a good choice for workloads that don’t use the full CPU often or consistently, but occasionally need to burst. Example workloads include virtual desktops, web servers, developer environments and small databases. For more information, see “Burstable Performance Instances” at <http://aws.amazon.com/ec2/instance-types/#burst>.

These instance types are not only suited for Hosted VDI Desktops but also for small XenApp workers, especially for the Task Worker workload, due to their ability to host a relatively large amount of XenApp Hosted Shared Desktop sessions for a very low cost per user (as low as \$0.006 – \$0.007 for the t2.small and t2.medium instances types). These instance types also

allow scaling up or down a XenApp or XenDesktop Delivery Group without impacting a large amount of users.

**Table 3. T2 instance specifications.**

Model	vCPU	Baseline Performance	CPU Credits/hour	Mem (GiB)	Storage (GiB)
t2.micro	1	10%	6	1	EBS Only
t2.small	1	20%	12	2	EBS Only
t2.medium	2	40%	24	4	EBS Only

### General Purpose T2 Small Instance (t2.small)

This section details the results from the XenApp Hosted Shared Desktop single instance validation testing on General Purpose T2 Small Instance with the API.name t2.small.

A General Purpose T2 Small instance comes with 2 GB of memory, 1 vCPU (1 core of an Intel Xeon E5-2670 v2, “Sandy Bridge” architecture), and EBS storage. It has been configured with the operating system and applications on the EBS boot volume, as well as the pagefile and user profiles, as the T2 instance types have no local instance storage.

Configured in this way, it allowed for six Hosted Shared Desktops with the Login VSI Task Worker workload.

The following graphs (Figure 16 to Figure 25) detail CPU, memory, disk and network performance on the Amazon EC2 t2.small instance. Memory was the limiting factor, as you can see from Figure 26. Available memory and % committed bytes in use. This result is understandable, given that this instance type has only 2 GB of memory.

Figure 16 contains the results for Task Worker desktop sessions, running Windows Server 2012R2 and Office 2013, and Figure 17 shows the corresponding CPU utilization.

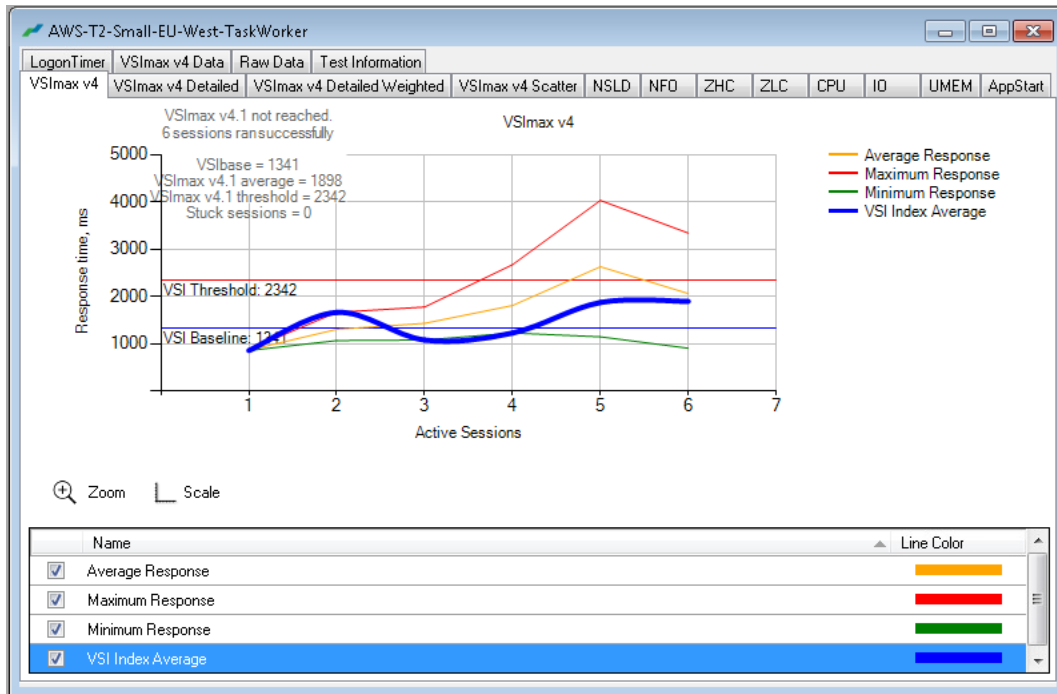


Figure 16. Task Worker desktop sessions on t2.small (Windows Server 2012R2 and Office 2013).

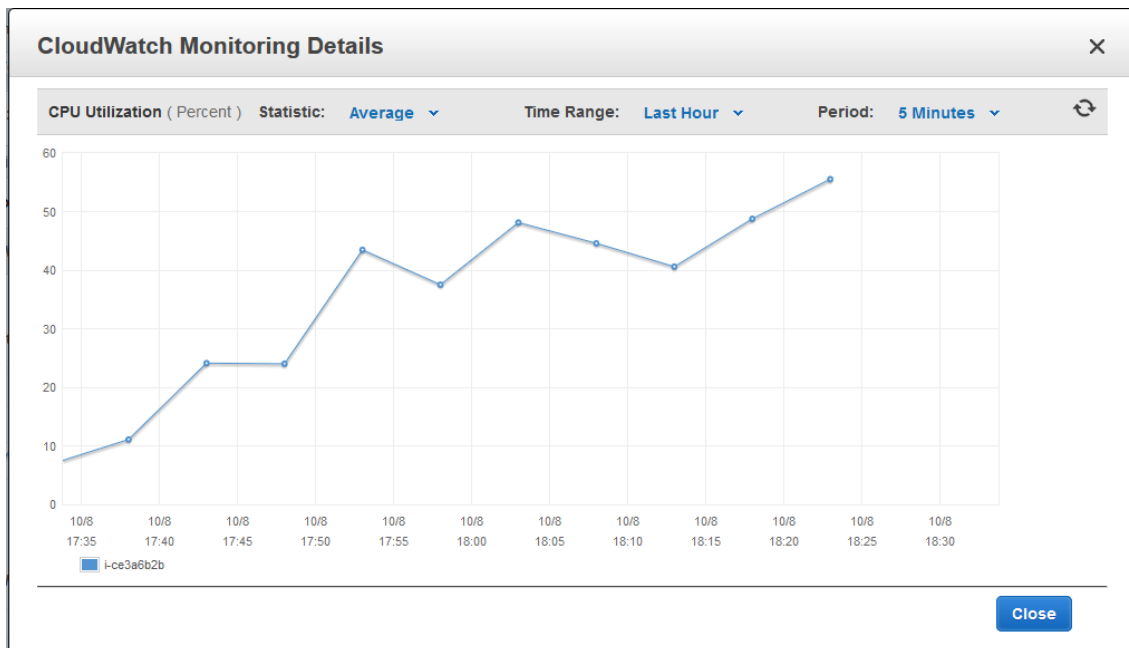


Figure 17. CPU Utilization—CloudWatch.

## CPU Credits

Now let's look at the CPU Credits while using the T2 instances with a desktop virtualization role such as XenDesktop and XenApp. CPU Credits accumulate when the instance doesn't use its baseline allocation of CPU, and they are spent when the instance is active. Unused CPU Credits are stored for up to 24 hours.

As listed in Table 3 on page 26, each T2 instance receives CPU Credits at a rate that is determined by the size of the instance. A CPU Credit provides the performance of a full CPU core for one minute.

In our example, a t2.small instance receives credits continuously at a rate of 12 CPU Credits per hour. This capability provides baseline performance equivalent to 20% of a CPU core. If at any moment the instance does not need the credits it receives, it stores them in its CPU Credit balance for up to 24 hours. If your instance doesn't use its baseline CPU for 5 hours, for example, the t2.small instance will have accumulated enough credits to run for almost an hour with full core performance (5 hours \* 12 CPU Credits / hour = 60 CPU Credits).

Desktops typically have bursts of activity (going through emails in the morning, launching applications) that are often followed by moments of inactivity (while users attend meetings, discuss information received by email). By putting a virtual desktop or application on a T2 instance, you can handle the compute load at active times expeditiously and cost effectively using the CPU Credits that were accumulated during the non-active times during the day.

As you can see in the graph of a 1-hour Login VSI Task Worker test run (Figure 18), we've only consumed 25 credits of our initial 80. Which would mean that after 3 hours of usage we would almost have depleted the amount of available credits, bringing us back to the 20% of one CPU core only. Although a t2.small instance for one hour would be able to support 6 Task Worker users, it is advised to not support more than 4 users. Depleting the CPU Credits in this way at a lower rate, and combining this with natural inactivity during coffee breaks, lunch breaks and a 1-2 hours "warmup" period (pre-starting instances 1-2 hours before working hours, accumulating 12 per hour bonus), prevents your t2.small instances from running out of CPU Credits and hence having a non-optimal user experience.




Figure 18. Task Worker desktop sessions on T2.small (Windows Server 2012R and Office 2013).

Figure 19 and Figure 20 show the CPU Credit usage and balance, respectively, for a t2.small instance with six Task Worker users.

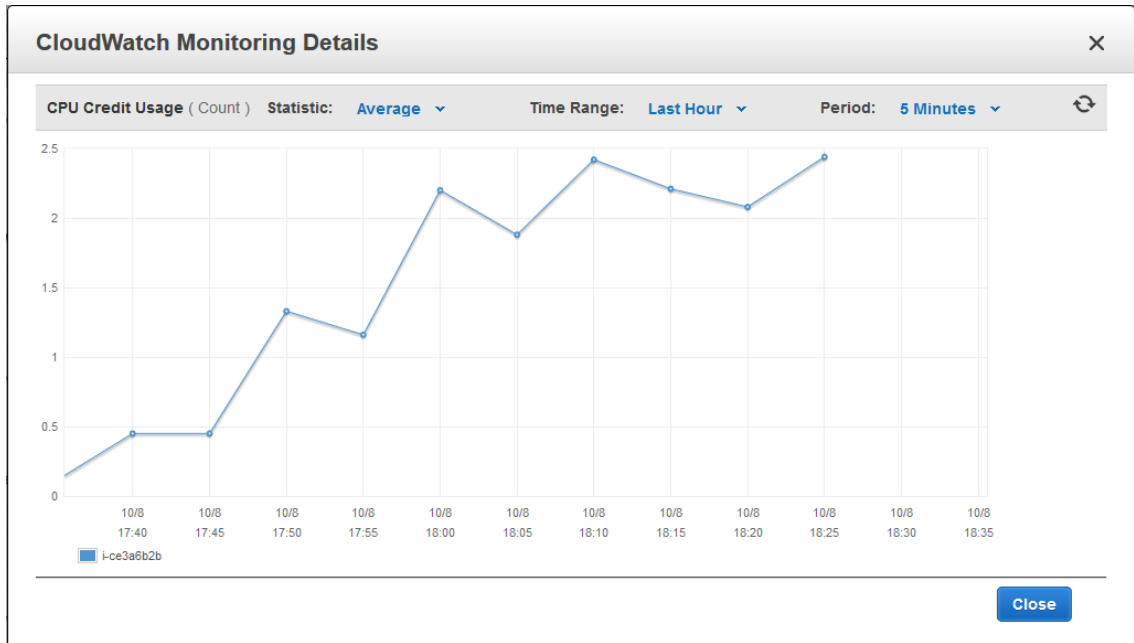


Figure 19. CPU Credit usage—t2.small—6 Task Worker users.




Figure 20. CPU Credit balance—t2.small—6 Task Worker users.

Figure 21 and Figure 22 show the CPU credit balance and usage, respectively for a t2.small instance with four Task Worker users.




Figure 21. CPU Credit balance—t2.small—4 Task Worker users—4 XA instances.




Figure 22. CPU Credit Usage—t2.small—4 Task Worker users—4 XA instances.

Figure 23 shows the CPU utilization and processor queue length; Figure 24 shows the percentage of pagefile usage.




Figure 23. CPU utilization and processor queue length.




Figure 24. Percentage of pagefile usage.

Figure 25 shows the disk queue length; Figure 26 shows the available memory and percentage of committed bytes in use.

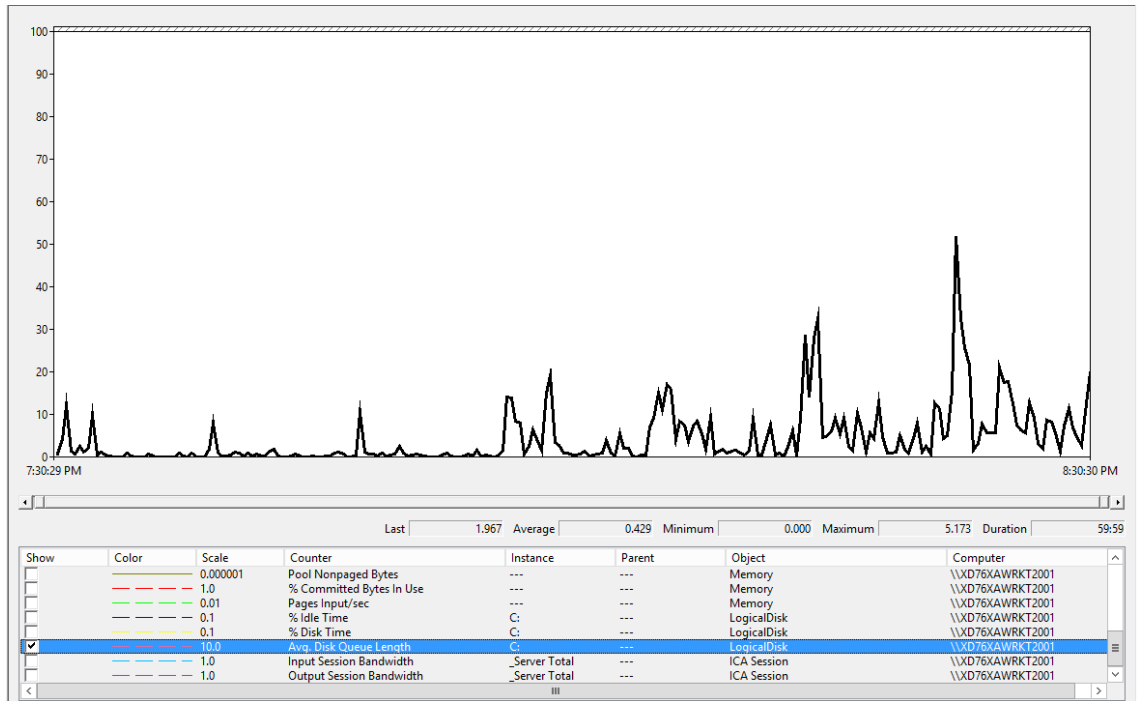


Figure 25. Disk queue length.




Figure 26. Available memory and % committed bytes in use.

## General Purpose T2 Medium Instance (t2.medium)

This section details the results from the XenApp Hosted Shared Desktop single instance validation testing on General Purpose T2 Medium Instance with the API.name t2.medium.

A General Purpose T2 Medium instance comes with 4 GB of memory, 2 vCPUs (2 cores of an Intel Xeon E5-2670 v2, “Sandy Bridge” architecture), and EBS storage. It has been configured



with the operating system and applications on the EBS boot volume, as well as the pagefile and user profiles, as the T2 instance types have no local instance storage.

Configured in this way, it allowed for 10 Hosted Shared Desktops with the Login VSI Task Worker workload. The following graphs detail CPU, memory, disk and network performance on the Amazon EC2 t2.medium instance. Memory utilization was the limiting factor, as you can see from the disk queue length increasing and the available memory dropping below 10% in Figure 30 and Figure 31.

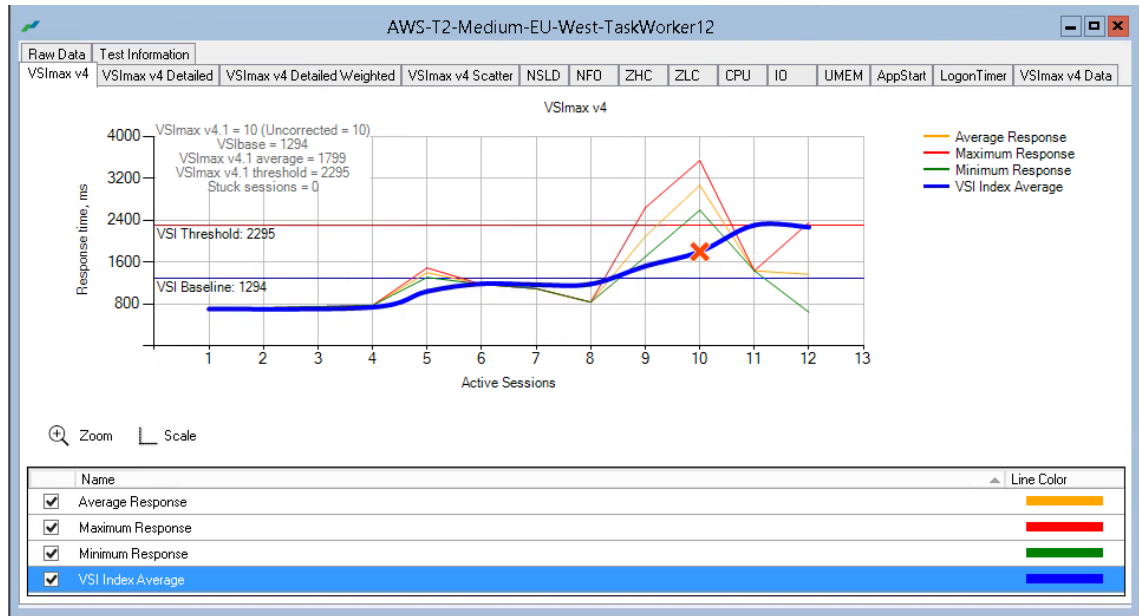


Figure 27. Task Worker Desktop Sessions on t2.medium Windows Server 2012 R2 Office 2013.

As you can see in the graph below of a 1-hour Login VSI Task Worker test run (Figure 28), we've only consumed 24 credits of our initial 61. Which would mean that after 3 hours of usage we would have depleted the amount of available credits, bringing us back to the 40% of one CPU core only. Although a t2.medium instance for one hour would be able to support 10 Task Worker users, it is advised to not support more than 6 users. Depleting the CPU Credits in this way at a lower rate, and combining this with natural inactivity during coffee breaks, lunch breaks and a 1-2 hours "warmup" period (pre-starting instances 1-2 hours before working hours, accumulating 12 per hour bonus), prevents your t2.medium instances from running dry on CPU Credits and hence providing a non-optimal user experience.




Figure 28. CPU Credit balance—t2.medium.

Figure 29 shows the CPU utilization and processor queue length for a t2.medium instance under test.




Figure 29. CPU utilization and processor queue length.

Figure 30 shows the disk queue length; Figure 31 shows the memory utilization and percentage of committed bytes in use for a t2.medium instance under test.




Figure 30. Disk queue length.

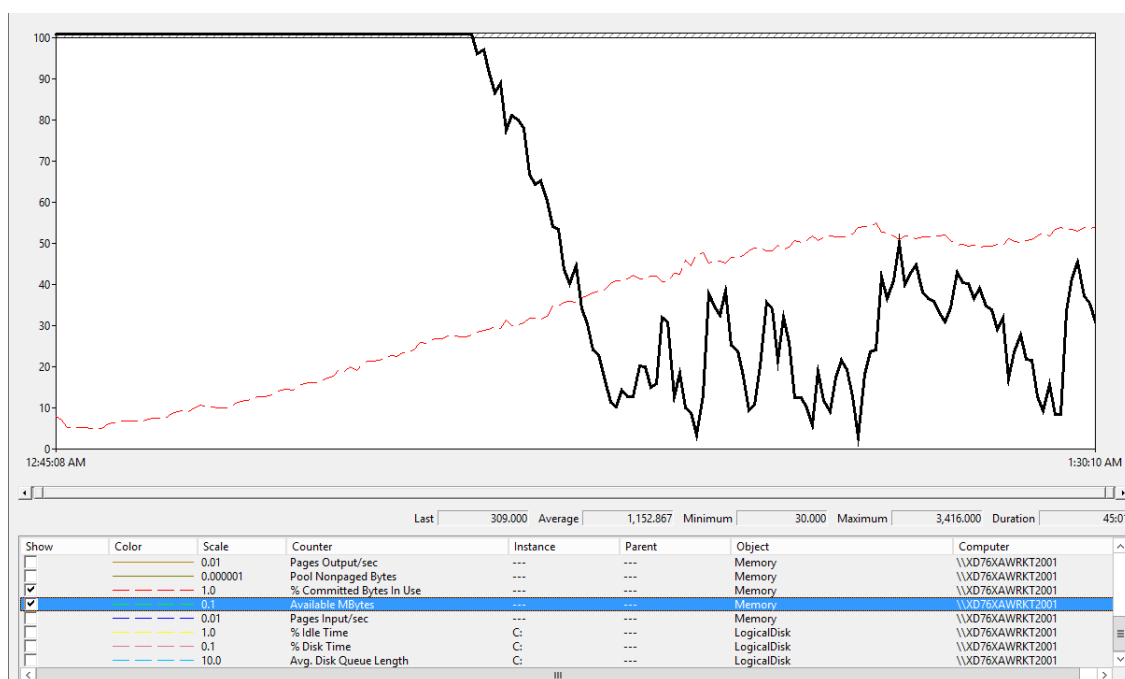


Figure 31. Memory utilization and % committed bytes in use.

## General Purpose M3 Double Extra Large Instance (m3.2xlarge)

This section details the results from the XenApp Hosted Shared Desktop single instance validation testing on General Purpose M3 Double Extra Large Instance with the API.name m3.2xlarge. Only these results are shown in this document; the respective scores for the other General Purpose M3 instance types can be found in Appendix B: Cost per user for different workloads overview on page 52.

Table 4. M3 instance specifications.

Model	vCPU	Mem (GiB)	SSD Storage (GB)
m3.medium	1	3.75	1 x 4
m3.large	2	7.5	1 x 32
m3.xlarge	4	15	2 x 40
m3.2xlarge	8	30	2 x 80

A General Purpose M3 Double Extra Large instance comes with 30 GB of memory, 8 vCPUs (2 cores of an Intel Xeon E5-2670 v2, “Ivy Bridge” architecture), EBS storage and 2 x 80 GB of local instance SSD storage. It has been configured with the operating system and applications on the EBS boot volume (magnetic EBS volume). The pagefile (drive letter Z: in the graphs) and user profiles (drive letter Y: in the graphs) are located on the local instance storage. This configuration offloads most of the volatile disk I/O to the fast local SSD instance storage.

Configured in this way, it allowed for 42 Hosted Shared Desktops with the Login VSI Task Worker workload. The following graphs (Figure 32–Figure 40) detail CPU, memory, and disk performance on the Amazon EC2 m3.2xlarge instance. Memory utilization was the limiting

factor, as you can see from the steep increase in paging activity to the Z: drive in Figure 36. Available memory versus paging disk IOPS m3.2xlarge.

Figure 32 shows response time statistics for Task Worker desktop sessions on an m3.2xlarge instance running Windows Server 2012R2 and Office 2013. Figure 33 shows a Citrix Director Monitor view of m3.2xlarge instance under test.

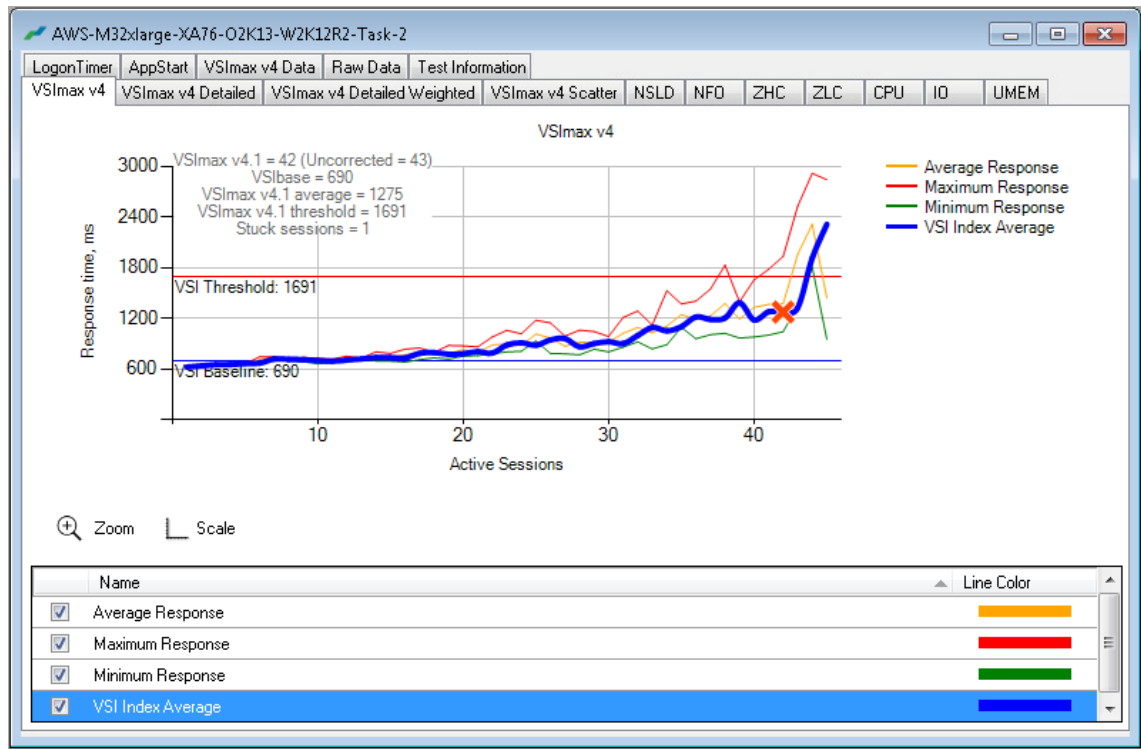


Figure 32. Task Worker desktop sessions on m3.2xlarge Windows Server 2012 R2, Office 2013.

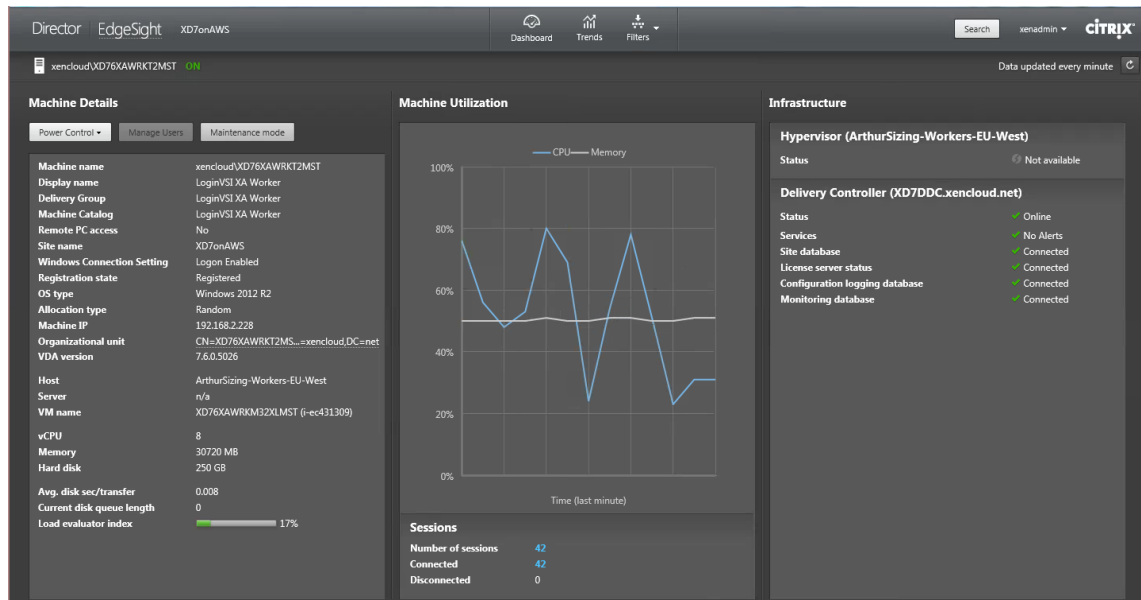


Figure 33. Citrix Director Monitor view of m3.2xlarge instance under test.

Figure 34 and Figure 35 show the inbound and outbound HDX traffic (in bps) for 42 sessions on an m3.2xlarge instance.




Figure 34. Inbound HDX traffic in bps for 42 sessions on m3.2xlarge.




Figure 35. Outbound HDX traffic in bps for 42 sessions on m3.2xlarge.

Figure 36 shows the available memory versus paging disk IOPS, and Figure 37 shows processor time versus active sessions, both for m3.2xlarge instances.




Figure 36. Available memory versus paging disk IOPS m3.2xlarge.




Figure 37. Processor time versus active sessions for m3.2xlarge.

Figure 38 and Figure 39 show the C drive IOPS and profile disk IOPS, respectively, versus the disk queue length for an m3.2xlarge instance.




Figure 38. C drive IOPS versus disk queue length for m3.2xlarge.




Figure 39. Profile disk IOPS versus disk queue length for m3.2xlarge.

Figure 40 shows the pagefile disk IOPS versus disk queue length for an m3.2xlarge instance.

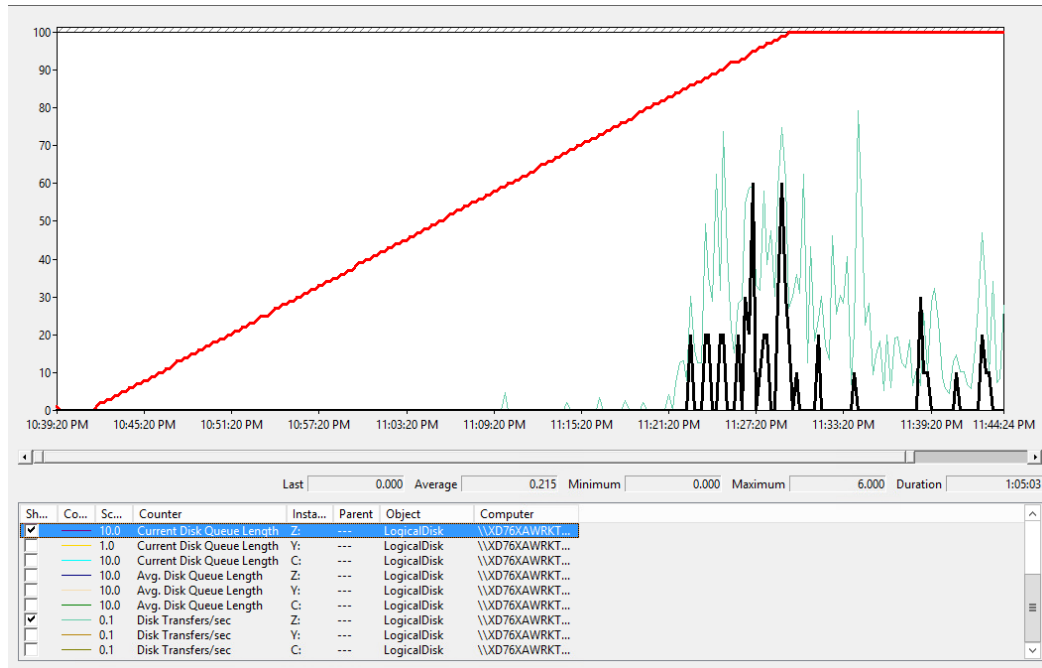


Figure 40. Pagefile disk IOPS versus disk queue length for m3.2xlarge.

## Compute Optimized C3 instances

C3 Instances are EC2's latest generation of compute-optimized instances, providing customers with the highest performing processors and a well-balanced CPU/Memory ratio to provide the second best price/user XenApp performance available in EC2 currently. See Appendix B: Cost per user for different workloads overview on page 52 for more details.

Table 5. C3 instance specifications.

Model	vCPU	Mem (GiB)	SSD Storage (GB)
c3.large	2	3.75	2 x 16
c3.xlarge	4	7.5	2 x 40
c3.2xlarge	8	15	2 x 80
c3.4xlarge	16	30	2 x 160
c3.8xlarge	32	60	2 x 320

## Compute Optimized C3 Quadruple Extra Large instance (c3.4xlarge)

This section details the results from the XenApp Hosted Shared Desktop single instance validation testing on Compute Optimized C3 Quadruple Extra Large Instance with the API.name c3.4xlarge. Only these results are shown in this document, the respective scores for the other General Purpose C3 instance types can be found in Appendix B: Cost per user for different workloads overview.

A Compute Optimized C3 Quadruple Extra Large instance comes with 30 GB of memory, 16 vCPUs (Intel Xeon E5-2680 v2, "Ivy Bridge" architecture), EBS storage and 2 x 160 GB of local instance SSD storage. It has been configured with the operating system and applications on the EBS boot volume (magnetic EBS volume), whereas the pagefile (drive letter Z: in the



graphs) and user profiles (drive letter Y: in the graphs) are located on the local instance storage. This configuration offloads most of the volatile disk I/O to the fast local SSD instance storage.

Configured in this way, it allowed for 83 Hosted Shared Desktops with the Login VSI Task Worker workload and 71 users with the Knowledge Worker workload. The following graphs detail CPU, memory, and disk performance on the Amazon EC2 c3.4xlarge instance. Memory utilization was the limiting factor, as can be seen by the percentage of committed bytes in use (see Figure 44). This is the ratio of committed bytes to the commit limit—in other words, the amount of virtual memory in use. A ratio greater than 80% indicates insufficient memory.

Figure 41 and Figure 42 show the Task Worker and Knowledge Worker desktop sessions, respectively.




Figure 41. Task Worker Desktop Sessions on c3.4xlarge Windows Server 2012 R2 Office 2013.

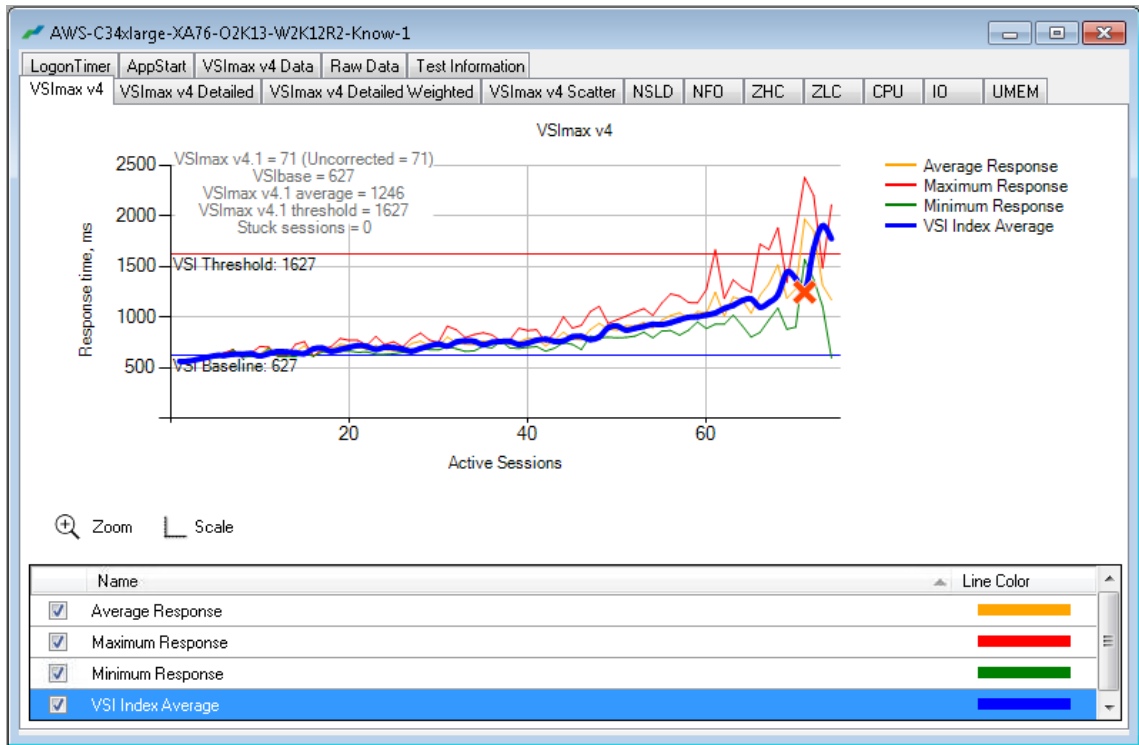


Figure 42. Knowledge Worker Desktop Sessions on c3.4xlarge Windows Server 2012 R2 Office 2013.

Figure 43 and Figure 44 show the Task Worker workload CPU performance and memory utilization, respectively.

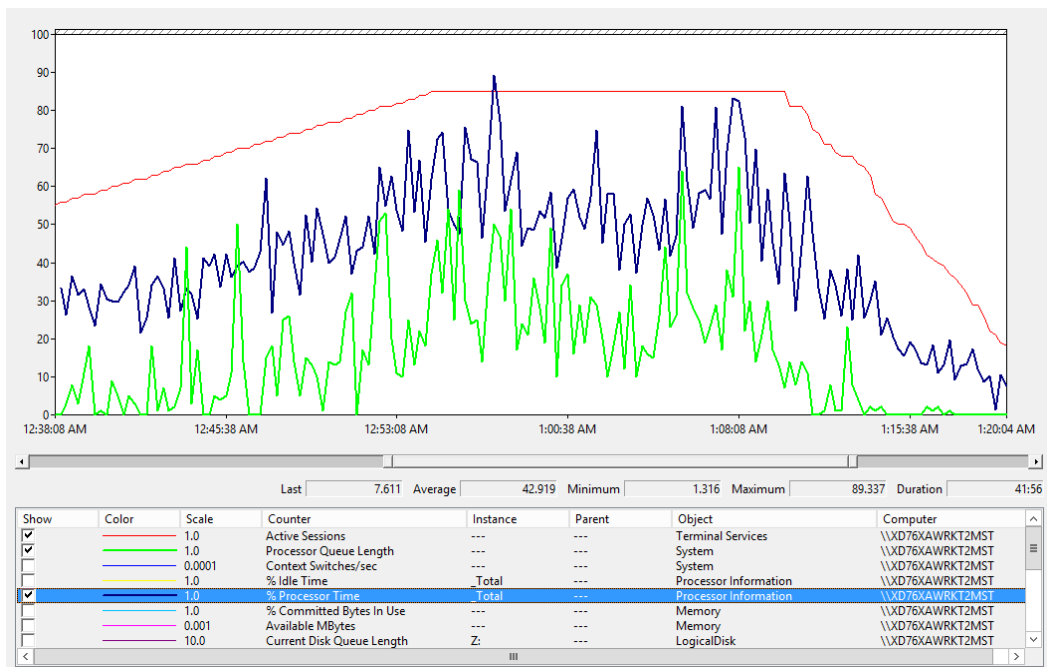


Figure 43. Task Worker workload CPU performance on c3.4xlarge in steady state for 84 session.

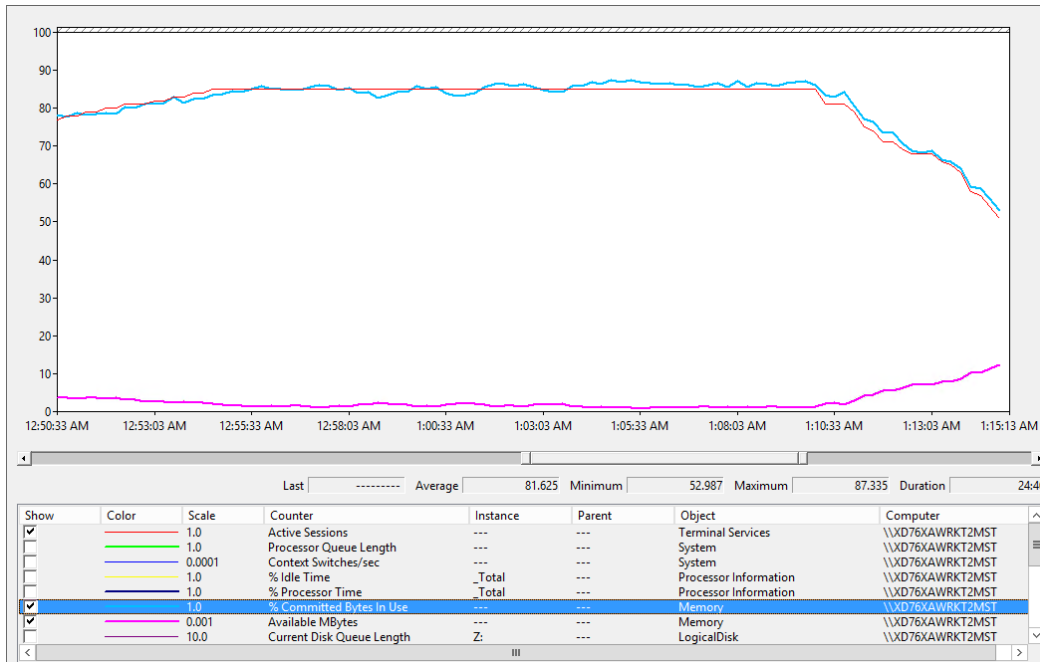


Figure 44. Task Worker workload memory utilization on c3.4xlarge in steady state for 84 session.

Figure 45 and Figure 46 show the C drive and profile drive IOPS, respectively, versus disk queue length for the c3.4xlarge Task Worker workload.




Figure 45. C Drive IOPS versus disk queue length for c3.4xlarge Task Worker workload.

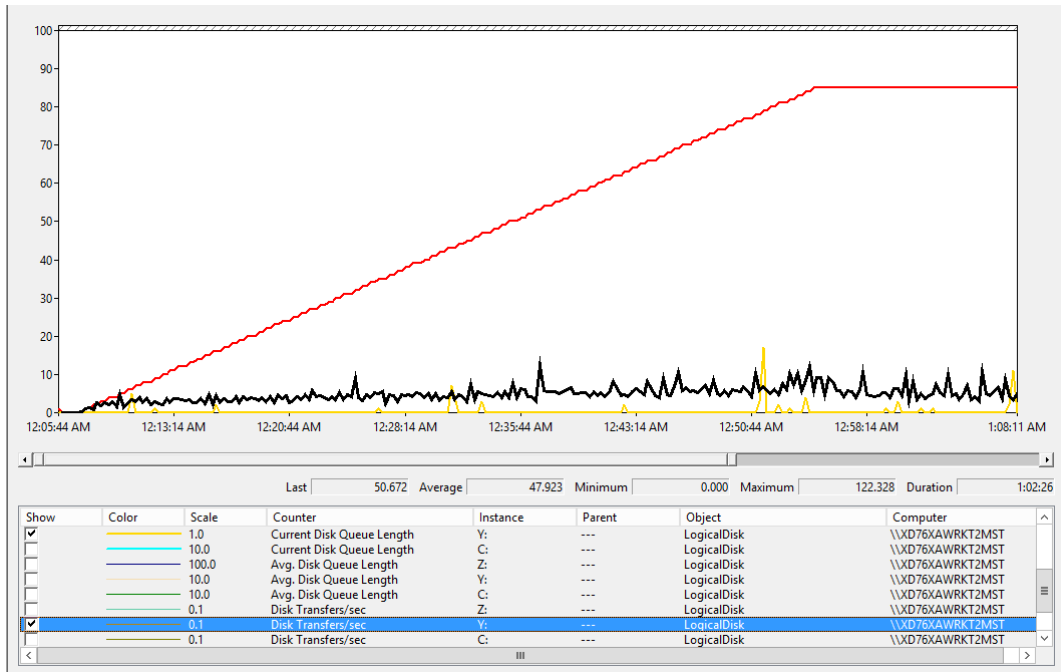


Figure 46. Profile disk IOPS versus disk queue length for c3.4xlarge.

## Test results: Multi-instance validation–1,000 Users

This section details the results from the XenApp/XenDesktop Hosted Shared Desktop sixteen instance validation testing. It illustrates linear scalability from one instance with 70 Users to 16 instances with 1,000 users.

Each instance was configured as a c3.4xlarge instance with a single EBS boot volume using the Login VSI Task Worker workload, since the purpose of this specific test was to demonstrate the feasibility of a large number of XenApp Worker Instances on AWS.

For this test, we used an additional Citrix internal instrument named STAT that is able to control the many Login VSI Launchers from a single central console as well as collect performance data from all involved components. This allows us to graph the progress of HDX sessions logging in successfully to the XenApp Farm.

Additional graphs detailing the CPU and utilization during peak session load are also presented. Given adequate storage capability, the limiting factor in the testing was CPU utilization. Each of the 16 c3.4xlarge instances performance charts are essentially identical for the multi-instance runs. We are including AWS CloudWatch data on all of the sixteen instances, for CPU, disk and network utilization.

1.000 Sessions							
Session Control ▾		Send Message					
<input type="checkbox"/>	Associated User	Session State	Session Start Time	Anonymous	Endpoint Name ▲	Endpoint IP	Receive
<input type="checkbox"/>	User000087	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039065	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000157	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000115	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039076	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000046	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039036	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039041	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039051	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000244	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User039145	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000106	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000118	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000143	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000202	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000179	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000097	Active	10/30/2014 5:34 PM	No	n/a	n/a	n/a
<input type="checkbox"/>	User000010	Active	10/30/2014 5:34 PM	No	simclient-10	207.198.47.68	6248.45
<input type="checkbox"/>	User010000	Active	10/30/2014 5:44 PM	No	simclient-10000	227.48.0.187	2474.22

Figure 47. 1000 Hosted Shared Desktop Session on 16 AWS EC2 c3.4xlarge instances.

Delivery Groups				
Delivery Group	Machine type	No. of machines	Sessions in use	No. of applications
PhyRDS		1	25	42
State: Enabled		Unregistered: 0	Disconnected: 6	
PhyWS		2	1	0
State: Enabled		Unregistered: 0	Disconnected: 0	
RDS		16	1000	0
State: Enabled		Unregistered: 0	Disconnected: 0	
Workstation		1000	0	0
State: Enabled		Unregistered: 0	Disconnected: 0	

Figure 48. 1000 Hosted Shared Desktop Sessions on 16 AWS EC2 c3.4xlarge instances.




Figure 49. 1000 Successful Hosted Shared Desktops on AWS.

Figure 50 and Figure 51 show CloudWatch metrics for 16 c3.4xlarge instances and all 16 EBS volumes, respectively.

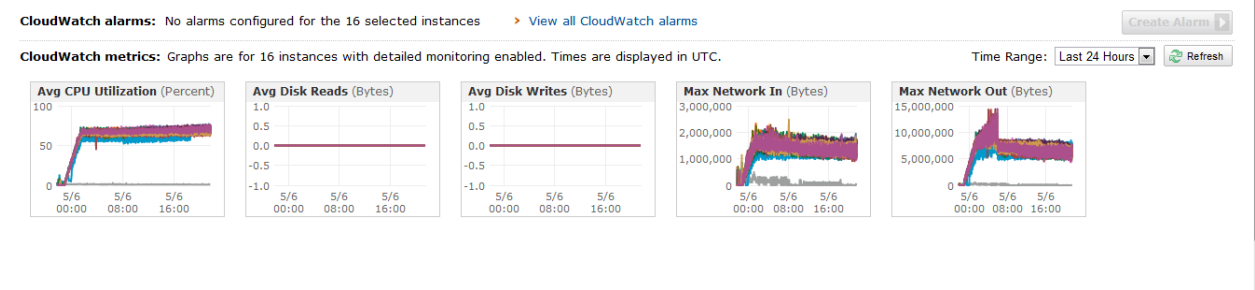


Figure 50. CloudWatch metrics for 16 c3.4xlarge instances.

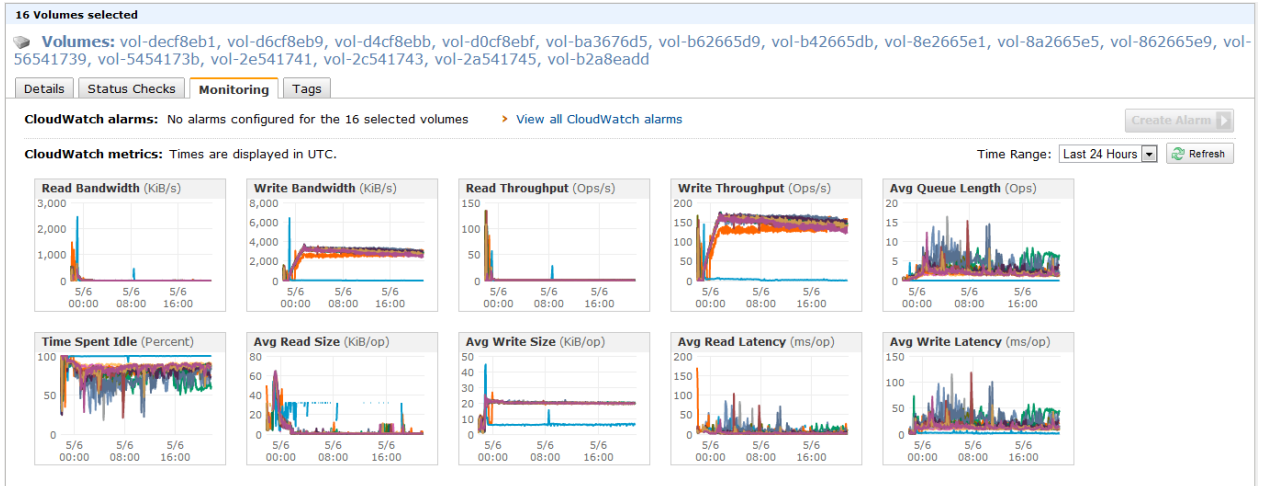


Figure 51. CloudWatch metrics for all 16 EBS volumes.

# Scalability considerations and guidelines

There are several factors to consider when you begin to deploy instances for a XenApp and XenDesktop 7.6 environment on AWS. In this section, we give some guidance on configuring the AWS instances.

## AWS Instance configuration

As our results indicate, we have proven linear scalability in the Amazon EC2 Reference Architecture as tested.

## AWS Storage best practices

Storage optimization for deploying XenApp Hosted Shared Desktop solutions on Amazon EBS includes the following:

- Use EBS root volumes to hold both operating system and XenApp installation.
- A fully loaded XenApp server instance can easily reach 200+ IOPS, which for a larger XenApp farm would become a very costly exercise for I/O operations related to volatile data—especially given that Amazon charges \$0.05 per million IOPS per month for magnetic disks. Several options exist:
  - Use general purpose SSDs instead of magnetic disks. There's no cost for associated IOPS, but storage allocation is \$0.10 per allocated GB per month versus the \$0.05 per allocated GB per month for the magnetic disk offering. General purpose SSD-based EBS volumes support bursts up to 3000 IOPS, and for a 50GB volume size guarantee a minimum of 150 IOPS.
  - Alternatively, use an instance type such as the c3.xlarge or c3.2xlarge that supports two local storage SSDs, which would contain:
    - Page file
    - User profiles
    - Application streaming cache

This would offload most of the I/O operations to the local SSDs which have no cost associated to them. They are, however, ephemeral, and any non-volatile data (such as user profiles) will have to be saved on shared storage upon logoff. A correct setup of user profiles is therefore a requirement.

Storage optimization for deploying XenDesktop VDI Desktop solutions on Amazon EBS is rather straight forward. As each single instance hosts only one single user, the amount of IOPS is limited to roughly 10 IOPS.

- Use EBS root volumes to hold the operating system.
- Use EBS magnetic disk volumes as the amount of IOPS generated is very low. Storage allocation is \$0.05 per allocated GB per month, and an additional \$0.05 per million IOPS per month. As a single user or desktop would generate around 5 million I/O operations per month, this would be just a few cents.

Here is an example of how to create an AWS instance to be used for XenApp servers with two local instance disks:

Create an Amazon Machine Image (AMI) with local storage enabled, which then would contain the pagefile and user profiles.



```
ec2-run-instances -region us-east- -block-device-mapping
"/dev/sdc=ephemeral10" -block-device-mapping "/dev/sdd=ephemeral11" ami-9d10d8f4
-g sg-36e2fb5a -t c3.2xlarge -k xencloudkey -z us-east-1e -s subnet-940547fc
```

This command will instantiate a new c3.2xlarge instance in VPC using a master AMI with a boot disk on EBS and two additional drives which are local storage.

- Drive Y: for Pagefile and App-V streaming cache
- Drive Z: for User Profiles and data

## How to change the UserProfile directory to point to drive E:

Change the RegKey for HKLM\Software\Microsoft\Windows NT\CurrentVersion\Profilelist\ProfilesDirectory to point to, for instance, Z:\Users instead of %SystemDrive%\Users.

Use this in combination with Microsoft Group Policy Objects (GPOs) to automatically delete cached roaming profiles after log off.

Because the local instance disks are volatile, after a stop and start of the instance, they will no longer contain any data. The EC2 Instance service will automatically initialize and format them.

## Appendix A: Performance counters

Table 6 contains recommendations and descriptions for performance counters.

**Table 6. Performance counters.**

Counter Name	Recommendation	Description
\Cache\Copy Read Hits %	≥ 98%	Percentage of cache copy read requests that hit the cache; that is, they did not require a disk read to provide access to the page in the cache. A copy read is a file read operation that is satisfied by a memory copy from a page in the cache to the application's buffer.
\LogicalDisk\% Disk Time	≤ 50%	Percentage of elapsed time that the selected disk is busy responding to read or write requests.
\LogicalDisk\% Idle Time	≥ 50%	Percentage of time during the sample interval that the disk was idle.
\LogicalDisk\Current Disk Queue Length	≤ 2	Number of requests outstanding on the disk at the time the performance data is collected.
\Memory\% Committed Bytes In Use	≤ 85%	Ratio of Memory\Committed Bytes to the Memory\Commit Limit. Committed memory is the physical memory in use for which space is reserved in the paging file if it needs to be written to disk.
\Memory\Pages Input/sec	≤ 100	Rate at which pages are read from disk to resolve hard page faults.
\Memory\Pages Output/sec	≤ 100	Rate at which pages are written to disk to free space in physical memory.
\Memory\Pool Nonpaged Bytes	≥ 200 MB	Size, in bytes, of the non-paged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated.
\Memory\Pool Paged Bytes	≥ 200 MB	Size, in bytes, of the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used.
\Network Interface\Bytes Received/sec	≤ Curr. Bandwidth	Rate at which bytes are received over each network adapter, including framing characters.
\Network Interface\Bytes Sent/sec	≤ Curr. Bandwidth	Rate at which bytes are sent over each network adapter, including framing characters.
\Network Interface\Bytes Total/sec	≤ Curr. Bandwidth	Rate at which bytes are sent and received over each network adapter, including framing characters.
\Network Interface\Current Bandwidth	NA	Estimate of the current bandwidth of the network interface in bits per second.
\Paging File\% Usage Peak	≥ 40%	Peak usage of the Page File instance in percent.
\Paging File\% Usage	≥ 40%	Amount of the Page File instance in use in percent.
\Processor\% Processor Time	≤ 65%	Percentage of elapsed time that the processor spends to execute a non-idle thread.

\System\Context Switches/sec	≤ 12000 per Core	The combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service.
\System\Processor Queue Length	≤ 2 per Core	Number of threads in the processor queue.
\Terminal Services\Active Sessions	NA	Actively connected terminal services sessions.

## Appendix B: Cost per user for different workloads overview

The following tables contain the cost estimates for the Task Worker, Knowledge Worker, and Power User workloads. Tables shown in orange contain information for Windows Server 2012R2 and Office 2013. Tables shown in green contain information for Windows Server 2008R2 and Office 2010.

**Table 7. Cost estimates for Task Worker workload on Windows Server 2012R2, Office 2013.**

Task Worker Workload - XenApp 7.6 - Windows Server 2012R2 - Office 2013							
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	Cost per User/h	
<b>General Purpose</b>							
t2.micro	Var	1	1	\$ 0.018	1	\$	0.018
t2.small	Var	2	1	\$ 0.036	6	\$	0.006
t2.medium	Var	4	2	\$ 0.072	10	\$	0.007
m3.medium	3	3.75	1	\$ 0.133	4	\$	0.033
m3.large	6.5	7.5	2	\$ 0.266	9	\$	0.030
m3.xlarge	13	15	4	\$ 0.532	19	\$	0.028
m3.2xlarge	26	30	8	\$ 1.064	42	\$	0.025
<b>Compute Optimized</b>							
c3.large	7	3.75	2	\$ 0.188	10	\$	0.0188
c3.xlarge	14	7.5	4	\$ 0.376	22	\$	0.0171
c3.2xlarge	28	15	8	\$ 0.752	42	\$	0.0179
c3.4xlarge	55	30	16	\$ 1.504	83	\$	0.0181
c3.8xlarge	108	60	32	\$ 3.008	164	\$	0.0226
<b>Memory Optimized</b>							
r3.large	6.5	15	2	\$ 0.300	9	\$	0.033
r3.xlarge	13	30.5	4	\$ 0.600	18	\$	0.033
r3.2xlarge	26	61	8	\$ 1.080	35	\$	0.031
r3.4xlarge	52	122	16	\$ 1.944	69	\$	0.028
r3.8xlarge	104	244	32	\$ 3.500	135	\$	0.026
<b>Storage Optimized</b>							
i2.xlarge	14	30.5	4	\$ 0.973	18	\$	0.054
i2.2xlarge	27	61	8	\$ 1.946	35	\$	0.056
i2.4xlarge	53	122	16	\$ 3.891	69	\$	0.056
i2.8xlarge	104	244	32	\$ 7.782	135	\$	0.058

Table 8. Cost estimates for Task Worker workload on Windows Server 2008R2, Office 2010.

Task Worker Workload - XenApp 7.6 - Windows Server 2008R2 - Office 2010							
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	Cost per User/h	
<b>General Purpose</b>							
t2.micro	Var	1	1	\$ 0.018	1	\$	0.018
t2.small	Var	2	1	\$ 0.036	6	\$	0.006
t2.medium	Var	4	2	\$ 0.072	10	\$	0.007
m3.medium	3	3.75	1	\$ 0.133	5	\$	0.027
m3.large	6.5	7.5	2	\$ 0.266	11	\$	0.024
m3.xlarge	13	15	4	\$ 0.532	22	\$	0.024
m3.2xlarge	26	30	8	\$ 1.064	44	\$	0.024
<b>Compute Optimized</b>							
c3.large	7	3.75	2	\$ 0.188	12	\$	0.0157
c3.xlarge	14	7.5	4	\$ 0.376	24	\$	0.0157
c3.2xlarge	28	15	8	\$ 0.752	47	\$	0.0160
c3.4xlarge	55	30	16	\$ 1.504	92	\$	0.0163
c3.8xlarge	108	60	32	\$ 3.008	180	\$	0.0178
<b>Memory Optimized</b>							
r3.large	6.5	15	2	\$ 0.300	11	\$	0.027
r3.xlarge	13	30.5	4	\$ 0.600	22	\$	0.027
r3.2xlarge	26	61	8	\$ 1.080	43	\$	0.025
r3.4xlarge	52	122	16	\$ 1.944	87	\$	0.022
r3.8xlarge	104	244	32	\$ 3.500	173	\$	0.020
<b>Storage Optimized</b>							
i2.xlarge	14	30.5	4	\$ 0.973	24	\$	0.041
i2.2xlarge	27	61	8	\$ 1.946	45	\$	0.043
i2.4xlarge	53	122	16	\$ 3.891	89	\$	0.044
i2.8xlarge	104	244	32	\$ 7.782	173	\$	0.045

Table 9. Cost estimates for Knowledge Worker workload on Windows Server 2012R2, Office 2013.

Knowledge Worker Workload - XenApp 7.6 - Windows Server 2012R2 - Office 2013						
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	per User
<b>General Purpose</b>						
t2.micro	Var	1	1	\$ 0.018	1	\$ 0.018
t2.small	Var	2	1	\$ 0.036	6	\$ 0.009
t2.medium	Var	4	2	\$ 0.072	10	\$ 0.009
m3.medium	3	3.75	1	\$ 0.133	4	\$ 0.044
m3.large	6.5	7.5	2	\$ 0.266	9	\$ 0.038
m3.xlarge	13	15	4	\$ 0.532	19	\$ 0.038
m3.2xlarge	26	30	8	\$ 1.064	42	\$ 0.038
<b>Compute Optimized</b>						
c3.large	7	3.75	2	\$ 0.188	10	\$ 0.0235
c3.xlarge	14	7.5	4	\$ 0.376	22	\$ 0.0198
c3.2xlarge	28	15	8	\$ 0.752	42	\$ 0.0209
c3.4xlarge	55	30	16	\$ 1.504	83	\$ 0.0212
c3.8xlarge	108	60	32	\$ 3.008	164	\$ 0.0215
<b>Memory Optimized</b>						
r3.large	6.5	15	2	\$ 0.300	9	\$ 0.038
r3.xlarge	13	30.5	4	\$ 0.600	18	\$ 0.038
r3.2xlarge	26	61	8	\$ 1.080	35	\$ 0.036
r3.4xlarge	52	122	16	\$ 1.944	69	\$ 0.034
r3.8xlarge	104	244	32	\$ 3.500	135	\$ 0.030
<b>Storage Optimized</b>						
i2.xlarge	14	30.5	4	\$ 0.973	18	\$ 0.061
i2.2xlarge	27	61	8	\$ 1.946	35	\$ 0.065
i2.4xlarge	53	122	16	\$ 3.891	69	\$ 0.067
i2.8xlarge	104	244	32	\$ 7.782	135	\$ 0.068

Table 10. Cost estimates for Knowledge Worker workload on Windows Server 2008R2, Office 2010.

Knowledge Worker Workload - XenApp 7.6 - Windows Server 2008R2 - Office 2010						
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	per User
<b>General Purpose</b>						
t2.micro	Var	1	1	\$ 0.018	1	\$ 0.018
t2.small	Var	2	1	\$ 0.036	4	\$ 0.009
t2.medium	Var	4	2	\$ 0.072	8	\$ 0.009
m3.medium	3	3.75	1	\$ 0.133	4	\$ 0.033
m3.large	6.5	7.5	2	\$ 0.266	9	\$ 0.030
m3.xlarge	13	15	4	\$ 0.532	18	\$ 0.030
m3.2xlarge	26	30	8	\$ 1.064	37	\$ 0.029
<b>Compute Optimized</b>						
c3.large	7	3.75	2	\$ 0.188	10	\$ 0.0188
c3.xlarge	14	7.5	4	\$ 0.376	20	\$ 0.0188
c3.2xlarge	28	15	8	\$ 0.752	39	\$ 0.0193
c3.4xlarge	55	30	16	\$ 1.504	78	\$ 0.0193
c3.8xlarge	108	60	32	\$ 3.008	154	\$ 0.0195
<b>Memory Optimized</b>						
r3.large	6.5	15	2	\$ 0.300	9	\$ 0.033
r3.xlarge	13	30.5	4	\$ 0.600	18	\$ 0.033
r3.2xlarge	26	61	8	\$ 1.080	37	\$ 0.029
r3.4xlarge	52	122	16	\$ 1.944	74	\$ 0.026
r3.8xlarge	104	244	32	\$ 3.500	115	\$ 0.030
<b>Storage Optimized</b>						
i2.xlarge	14	30.5	4	\$ 0.973	15	\$ 0.065
i2.2xlarge	27	61	8	\$ 1.946	30	\$ 0.065
i2.4xlarge	53	122	16	\$ 3.891	59	\$ 0.066
i2.8xlarge	104	244	32	\$ 7.782	115	\$ 0.068

Table 11. Cost estimates for Power Worker workload on Windows Server 2012R2, Office 2013.

Power User Workload - XenApp 7.6 - Windows Server 2012R2 - Office 2013						
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	Cost per User/h
<b>General Purpose</b>						
t2.micro	Var	1	1	\$ 0.018	0	N/A
t2.small	Var	2	1	\$ 0.036	2	\$ 0.018
t2.medium	Var	4	2	\$ 0.072	6	\$ 0.012
m3.medium	3	3.75	1	\$ 0.133	1	\$ 0.133
m3.large	6.5	7.5	2	\$ 0.266	4	\$ 0.067
m3.xlarge	13	15	4	\$ 0.532	9	\$ 0.059
m3.2xlarge	26	30	8	\$ 1.064	20	\$ 0.053
<b>Compute Optimized</b>						
c3.large	7	3.75	2	\$ 0.188	7	\$ 0.0269
c3.xlarge	14	7.5	4	\$ 0.376	15	\$ 0.0251
c3.2xlarge	28	15	8	\$ 0.752	29	\$ 0.0259
c3.4xlarge	55	30	16	\$ 1.504	60	\$ 0.0251
c3.8xlarge	108	60	32	\$ 3.008	119	\$ 0.0253
<b>Memory Optimized</b>						
r3.large	6.5	15	2	\$ 0.300	5	\$ 0.060
r3.xlarge	13	30.5	4	\$ 0.600	11	\$ 0.055
r3.2xlarge	26	61	8	\$ 1.080	21	\$ 0.051
r3.4xlarge	52	122	16	\$ 1.944	43	\$ 0.045
r3.8xlarge	104	244	32	\$ 3.500	85	\$ 0.041
<b>Storage Optimized</b>						
i2.xlarge	14	30.5	4	\$ 0.973	11	\$ 0.088
i2.2xlarge	27	61	8	\$ 1.946	21	\$ 0.093
i2.4xlarge	53	122	16	\$ 3.891	43	\$ 0.090
i2.8xlarge	104	244	32	\$ 7.782	85	\$ 0.092



Table 12. Cost estimates for Power Worker workload on Windows Server 2008R2, Office 2010.

Power User Workload - XenApp 7.6 - Windows Server 2008R2 - Office 2010						
Instance type	Units	(GiB)	vCPUs	Cost/Hour	Sessions	Cost per User/h
<b>General Purpose</b>						
t2.micro	Var	1	1	\$ 0.018	0	N/A
t2.small	Var	2	1	\$ 0.036	2	\$ 0.018
t2.medium	Var	4	2	\$ 0.072	6	\$ 0.012
m3.medium	3	3.75	1	\$ 0.133	3	\$ 0.044
m3.large	6.5	7.5	2	\$ 0.266	7	\$ 0.038
m3.xlarge	13	15	4	\$ 0.532	14	\$ 0.038
m3.2xlarge	26	30	8	\$ 1.064	28	\$ 0.038
<b>Compute Optimized</b>						
c3.large	7	3.75	2	\$ 0.188	8	\$ 0.0235
c3.xlarge	14	7.5	4	\$ 0.376	17	\$ 0.0221
c3.2xlarge	28	15	8	\$ 0.752	33	\$ 0.0228
c3.4xlarge	55	30	16	\$ 1.504	67	\$ 0.0224
c3.8xlarge	108	60	32	\$ 3.008	130	\$ 0.0231
<b>Memory Optimized</b>						
r3.large	6.5	15	2	\$ 0.300	7	\$ 0.043
r3.xlarge	13	30.5	4	\$ 0.600	14	\$ 0.043
r3.2xlarge	26	61	8	\$ 1.080	28	\$ 0.039
r3.4xlarge	52	122	16	\$ 1.944	57	\$ 0.034
r3.8xlarge	104	244	32	\$ 3.500	94	\$ 0.037
<b>Storage Optimized</b>						
i2.xlarge	14	30.5	4	\$ 0.973	12	\$ 0.081
i2.2xlarge	27	61	8	\$ 1.946	24	\$ 0.081
i2.4xlarge	53	122	16	\$ 3.891	48	\$ 0.081
i2.8xlarge	104	244	32	\$ 7.782	94	\$ 0.083

## Appendix C: Online resources

### Amazon reference documents

Active Directory Reference Architecture

<http://aws.amazon.com/microsoft/whitepapers/ad-reference-architecture/>

SQL Server High Availability

<http://aws.amazon.com/microsoft/whitepapers/microsoft-wsfc-sql-alwayson/>

Secure Microsoft Applications on AWS

<http://aws.amazon.com/microsoft/whitepapers/secure-microsoft-applications-on-aws/>

### Citrix reference documents

Citrix XenDesktop with Login Consultants VSI

<http://support.citrix.com/article/CTX125630>

Login Consultants VSI References

<http://www.loginvsi.com/en/in-practice/references>

Citrix Virtual Desktop Handbook 7.x

<http://support.citrix.com/article/CTX139331>

How to Deploy XenApp and XenDesktop with Amazon Web Services

<http://support.citrix.com/article/CTX140427>

XenDesktop 7.5 and XenApp 7.5 working in Hybrid harmony with AWS

<http://blogs.citrix.com/2014/03/28/xendesktop-7-5-and-xenapp-7-5-working-in-hybrid-harmony-with-aws/>

XenDesktop 7.5 on AWS Design Guide

[http://www.citrix.com/content/dam/citrix/en\\_us/documents/partner-documents/xendesktop-75-on-amazon-web-services-design-guide.pdf](http://www.citrix.com/content/dam/citrix/en_us/documents/partner-documents/xendesktop-75-on-amazon-web-services-design-guide.pdf)

Scalability and Economics of XenApp 6.5 on Amazon cloud

[http://www.citrix.com/content/dam/citrix/en\\_us/documents/partner-documents/scalability-and-economics-of-xenapp-on-amazon-cloud.pdf](http://www.citrix.com/content/dam/citrix/en_us/documents/partner-documents/scalability-and-economics-of-xenapp-on-amazon-cloud.pdf)

Citrix XenApp 6.5 on AWS: Reference Architecture

[http://www.citrix.com/content/dam/citrix/en\\_us/documents/products-solutions/citrix-xenapp-on-aws-reference-architecture.pdf](http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/citrix-xenapp-on-aws-reference-architecture.pdf)

Citrix XenApp 6.5 on AWS: Implementation Guide

[http://www.citrix.com/content/dam/citrix/en\\_us/documents/products-solutions/citrix-xenapp-on-aws-implementation-guide.pdf](http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/citrix-xenapp-on-aws-implementation-guide.pdf)

**Corporate Headquarters**

Fort Lauderdale, FL, USA

**Silicon Valley Headquarters**

Santa Clara, CA, USA

**EMEA Headquarters**

Schaffhausen, Switzerland

**India Development Center**

Bangalore, India

**Online Division Headquarters**

Santa Barbara, CA, USA

**Pacific Headquarters**

Hong Kong, China

**Latin America Headquarters**

Coral Gables, FL, USA

**UK Development Center**

Chalfont, United Kingdom

**About Citrix**

Citrix (NASDAQ:CTXS) is a leader in mobile workspaces, providing virtualization, mobility management, networking and cloud services to enable new ways to work better. Citrix solutions power business mobility through secure, personal workspaces that provide people with instant access to apps, desktops, data and communications on any device, over any network and cloud. This year Citrix is celebrating 25 years of innovation, making IT simpler and people more productive. With annual revenue in 2013 of \$2.9 billion, Citrix solutions are in use at more than 330,000 organizations and by over 100 million users globally. Learn more at [www.citrix.com](http://www.citrix.com).

Copyright © 2014 Citrix Systems, Inc. All rights reserved. [List Citrix trademarks (without ® or ™ symbols!) in document] are trademarks of Citrix Systems, Inc. and/or one of its subsidiaries, and may be registered in the U.S. and other countries. Other product and company names mentioned herein may be trademarks of their respective companies.