



EMC Documentum xCP on VMware vSphere™ Virtual Infrastructure

December 2010

DEPLOYMENT AND TECHNICAL CONSIDERATIONS GUIDE

Table of Contents

Introduction.....	1
VMware and EMC Documentum xCelerated Composition Platform (xCP) Overview.....	1
EMC Documentum xCP	1
VMware vSphere or VMware Virtual Infrastructure	2
Documentum xCP Architecture and Deployment	3
Testing Process and Results.....	5
Testing Methodology and Overview	5
Observed Results.....	9
Business Continuity and High Availability Testing	13
High Availability.....	13
VMotion.....	14
Deployment Best Practices	15
Scaling up to Larger Configurations	16
Summary and Conclusions.....	17
Additional Resources.....	18

Introduction

This document provides direction to those interested in running EMC Documentum xCP on VMware Infrastructure 3 or VMware vSphere™ 4. It provides basic information on the architecture of Documentum xCP, as well as showing the value of utilizing the VMware platform for Documentum xCP deployments. The results of recent testing performed jointly by VMware and EMC are also included in this document and characterize the performance and functionality of Documentum xCP running on VMware virtual infrastructure. Finally, the document outlines some best practices for utilizing the two product sets together in your datacenter.

VMware and EMC Documentum xCelerated Composition Platform (xCP) Overview

VMware (NYSE: VMW), the global leader in virtualization and cloud infrastructure, delivers customer-proven solutions that significantly reduce IT complexity and enable more flexible, agile service delivery. VMware accelerates an organization's transition to cloud computing, while preserving existing IT investments. With more than 190,000 customers and 25,000 partners, VMware helps organizations of all sizes lower costs, preserve freedom of choice and energize business through IT, while saving energy—financial, human, and the Earth's.

EMC Documentum xCP

The EMC Documentum xCelerated Composition Platform (xCP) provides a better, faster way to build case management applications that manage your content-centric business processes. Case management automates structured and unstructured processes to improve performance, eliminate errors, and help organizations get to the right decisions.

With Documentum xCP, you can build case-based applications substantially faster, at a much lower cost, and with fewer resources – and reduce total cost of ownership by delivering applications that are easier to modify and maintain. Documentum xCP provides integrated development and deployment tools emphasizing configuration versus coding, uniting all the elements of case management into a single powerful environment. Documentum xCP unites content management, business process management, intelligent capture, customer communications management, collaboration case management, analysis, reporting and compliance into a single environment.

With Documentum xCP, system integrators (SIs), independent software vendors (ISVs), and IT organizations can also rapidly build intelligent case-based applications substantially faster, at lower cost, and with fewer resources.

VMware vSphere or VMware Virtual Infrastructure

VMware's leading virtualization solutions provide multiple benefits to IT administrators and users. VMware virtualization creates a layer of abstraction between the resources required by an application and operating system, and the underlying hardware that provides those resources. A summary of the value of this abstraction layer includes the following:

- **Consolidation:** VMware technology allows multiple application servers to be consolidated onto one physical server, with little or no decrease in overall performance.
- **Ease of Provisioning:** VMware virtualization encapsulates an application into an image that can be duplicated or moved, greatly reducing the cost of application provisioning and deployment.
- **Manageability:** Virtual machines may be moved from server to server with no downtime using VMware VMotion™, which simplifies common operations like hardware maintenance and reduces planned downtime.
- **Availability:** Unplanned downtime can be reduced and higher service levels can be provided to an application. VMware High Availability (HA) ensures that in the case of an unplanned hardware failure, any affected virtual machines are restarted on another host in a VMware cluster.

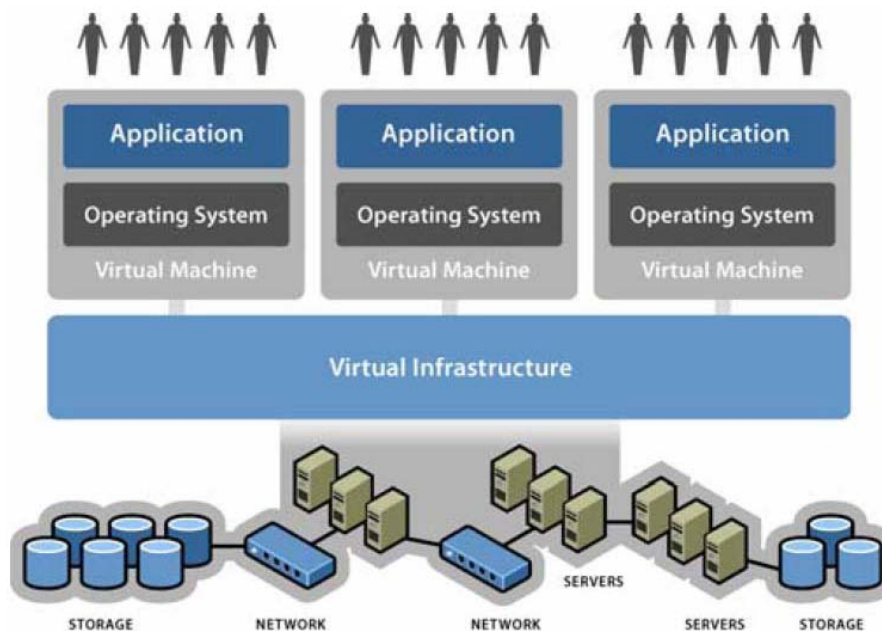


Figure 1. VMware vSphere Virtual Infrastructure

Documentum xCP Architecture and Deployment

A Documentum xCP system consists of the components shown in the following figure.

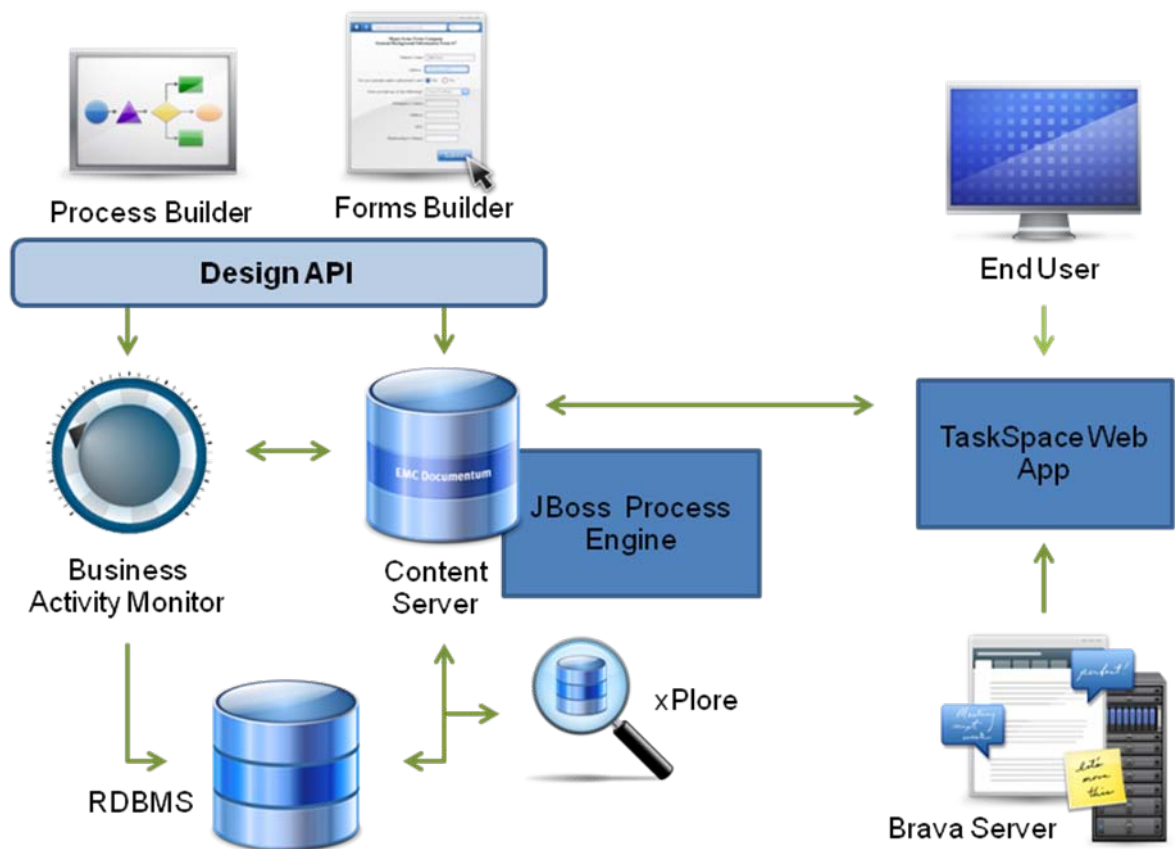


Figure 2. Documentum xCP Software Architecture

Content Server: The Documentum Content Server and repository manages information and content with enterprise-grade security, scalability, and reliability, and a common set of policies for all content assets across an organization. Additionally, the Documentum repository manages the application definition artifacts, configurations, and supporting models.

TaskSpace: Documentum TaskSpace provides a highly configurable user interface that unites process, content, and monitoring into a single user experience for transactional business applications. TaskSpace is deployed as a web application running with a Java Application server.

xPlore: Documentum xPlore provides powerful new search services for the Documentum platform. xPlore delivers superior search performance and scalability, improved supportability, and ease of use, and delivers a search and indexing capability fully supported in a virtual environment.

Process Builder: Process Builder is a business process design tool that allows non-programmers to configure how services, manual activities, and information are combined and orchestrated to accomplish specific business actions and logic. Process Builder allows the designer to configure process activities such as simple manual tasks, high-volume work queues, automated content management actions, and sophisticated system integration logic through easy-to-use visual design techniques. In addition, Process Builder supports the defining of process rules, priorities, and exception handling. Deployment of standard activity templates enables enterprises to maximize both productivity and standardization through reuse across multiple applications. The process models defined by Process Builder are deployed into Documentum repositories through the Documentum Content Server. The runtime execution of the process models is carried out by the Process Engine.

Forms Builder: Forms Builder is used to create web-based graphical user interfaces available through TaskSpace. Forms Builder also serves as the design tool for producing high-fidelity, paper-like electronic forms for data capture and presentation. Forms Builder produces XForms models that are stored in the Documentum Repository. At runtime, web-based user interfaces are generated by an embedded forms engine that interprets the XForms models and produces HTML. The forms engine deploys as part of TaskSpace in the same web application.

Business Activity Monitor (BAM): BAM provides detailed reporting and monitoring of solutions and is deployed in two parts, as a server-side engine (BAM Server) and as a TaskSpace component. The BAM Server continuously formats and aggregates events and data as the application runs. This engine deploys into any Java application server.

Process Engine: The Documentum Process Engine carries out the automation of all business processes defined by Process Builder in the applications. Process Engine deploys into the embedded JBoss application server that is provided as part of the Documentum Server. Certain aspects of the Process Engine also execute in the Documentum Content Server, specifically process runtime state management, which uses the Documentum Repository for state persistence.

BRAVA Server: The Brava! Server is a servlet-based application that responds to Brava! Client requests to access files and to retrieve and save markups. The Brava! Client which is installed with TaskSpace provides the GUI for viewing and marking up files.

For additional information on the deployment of each component, please refer to Documentum product installation guides.

Testing Process and Results

To characterize the performance of Documentum xCP on VMware virtual infrastructure, load tests at different levels were carried out jointly by VMware and EMC. The configuration tested and the results are summarized below.

Using LoadRunner to emulate a typical xCP application, we ran test cases at the following user and transaction rate levels:

- 50 users, approximately 2.5 transactions / second
- 100 users, approximately 3.0 transactions / second
- 200 users, approximately 6.0 transactions / second
- 50 “heavy” users with increased transaction rate that simulates a high load scenario that doubles the 200 user case, approximately 11.0 transactions / second

Table 1. Documentum xCP Application Test Cases

TEST CASE	RESPONSE TIME INDEX (SEC)	CPU UTILIZATION <TASKSPACE>	MEMORY UTILIZATION <TASKSPACE>	CPU UTILIZATION <CONTENT SERVER>	MEMORY UTILIZATION <CONTENT SERVER>	CPU UTILIZATION <DB>	MEMORY UTILIZATION <DB>
50 users (2.5 txn/sec)	0.35	15.5%	1650MB	5.9%	962MB	4.7%	746MB
100 users (3.0 txn/sec)	0.44	15.7%	2077MB	15.9%	1274MB	9.6%	647MB
200 users (6.0 txn/sec)	0.51	30.1%	2238MB	16.7%	1595MB	12.4%	630MB
Heavy users (11.0 txn/sec)	0.63	51.9%	2818MB	34.0%	2038MB	25.0%	914MB

Testing Methodology and Overview

The primary objectives of testing were to determine Documentum xCP performance characteristics and show that xCP can scale vertically and horizontally with increased load and be run more resiliently on VMware virtual infrastructure. Testing also validated the operation of VMware features including vMotion, VMware DRS, and VMware HA with xCP deployments in a virtual environment.

Hardware and Software Configuration

The following diagram provides details about the setup of ESX host servers and virtual machines used to perform the tests described earlier.

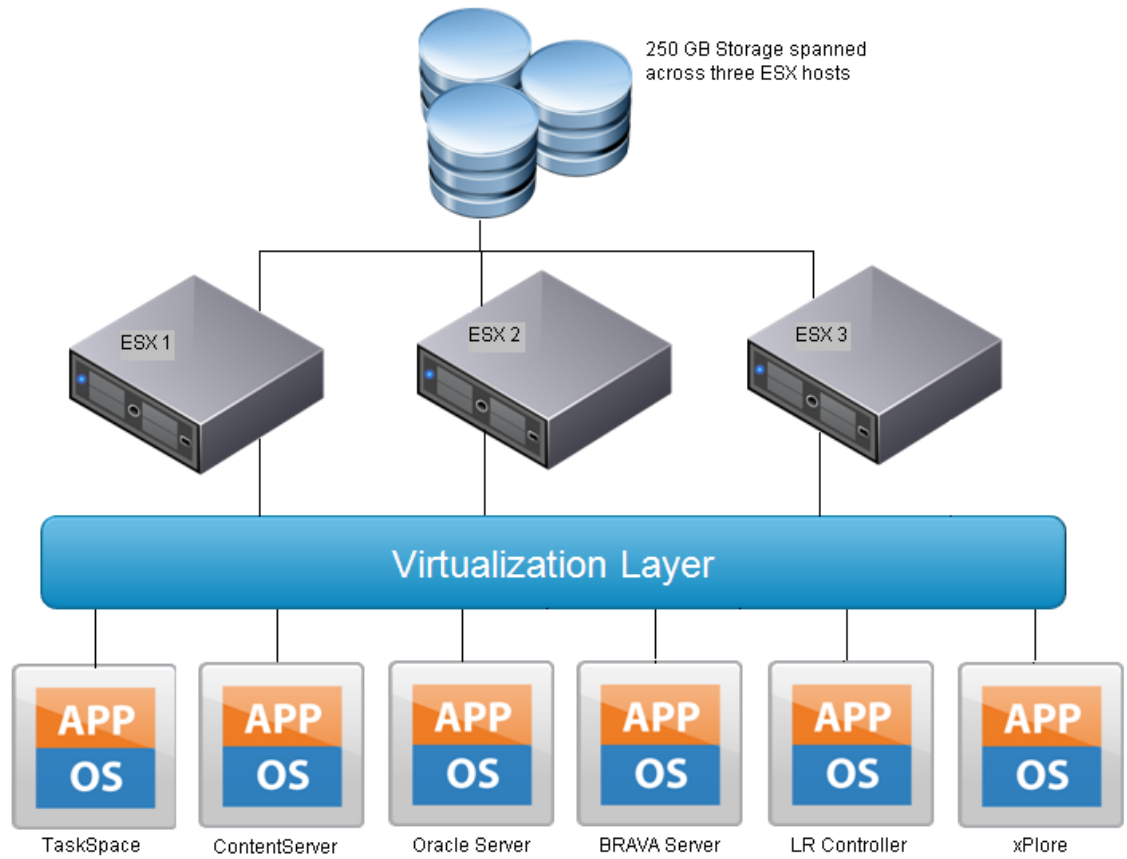


Figure 3. Physical ESX Host and Virtual Machine Configuration for Documentum xCP Testing

Hardware and System Host Configuration

The following table describes the configuration of ESX host servers and storage in the EMC Documentum test configurations.

Table 2. ESX Host Hardware

HARDWARE	CONFIGURATION
Server	Three HP Proliant BL 460c G6 (Nehalem) Servers. Each server is equipped with: <ul style="list-style-type: none"> • 2.8 GHz dual quad-core Intel Xeon E5560 • 48 GB RAM One HP Proliant BL 460c G1 Server. Each server is equipped with: <ul style="list-style-type: none"> • 3.0 GHz dual quad-core Intel Xeon E5450 • 48 GB RAM
Storage	EMC: <ul style="list-style-type: none"> • Clarion CX3-10, RAID 10 • FC 4Gbps • 15 disks, 133 GB each

Installed Software

Table 2 lists the software used in the EMC Documentum solution.

Table 3. Software Installed for EMC Documentum Solution

INSTALLED SOFTWARE	
VMware	VMware vSphere 4.1, vCenter™ Server 4.1, ESX® 4.1
xCP	Content Server 6.6 Process Engine 6.6 TaskSpace 6.6 xPlore 1.0 BAM 6.6 Brava Server 6.0 Process Builder 6.0 Forms Builder 6.6
Oracle	Oracle 11g 11.1.0.7
Microsoft	Microsoft Windows 2003 R2 Enterprise Edition

Virtual Machine Configuration

The following table describes the configuration of virtual machines running on ESX host servers in the xCP test configurations.

Table 4. xCP Virtual Machine Configurations

VIRTUAL MACHINE	HARDWARE CONFIGURATION	SOFTWARE INSTALLED
One TaskSpace virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 4 vCPUs • 4 GB memory • 1 Ethernet card • 1 X 20 GB • vmdk only 	<ul style="list-style-type: none"> • TaskSpace
One Content Server virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 4 vCPU • 4 GB memory • 1 Ethernet card • 1 X 80 GB • vmdk only 	<ul style="list-style-type: none"> • Content Server • Process Engine
One Oracle virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 4 vCPU • 4 GB memory • 1 Ethernet card • 1 X 100 GB • vmdk only 	<ul style="list-style-type: none"> • Oracle DB
One xPlore virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 4 vCPU • 4 GB memory • 1 Ethernet card • 1 X 100 GB • vmdk only 	<ul style="list-style-type: none"> • xPlore
One Brava Server virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 1 vCPU • 1 GB memory • 1 Ethernet card • 1 X 20 GB • vmdk only 	<ul style="list-style-type: none"> • Brava Server • Process Builder • Forms Builder
One Load Runner Controller machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 2 vCPU • 2 GB memory • 1 Ethernet card • 1 X 27 GB • vmdk only 	<ul style="list-style-type: none"> • LoadRunner

For the 200-User test, another TaskSpace virtual machine was deployed with the following configuration:

Table 5. Virtual Machine Configuration for 200-User Test

VIRTUAL MACHINE	HARDWARE CONFIGURATION	SOFTWARE INSTALLED
One TaskSpace virtual machine (Microsoft Windows 2003 R2 Enterprise Edition)	<ul style="list-style-type: none"> • 4 vCPUs • 4 GB memory • 1 Ethernet card • 1 X 20 GB • vmdk only 	<ul style="list-style-type: none"> • TaskSpace

Workload Used

For xCP performance testing, we employ a sample application constructed to model a loan approval process and its corresponding workflow. The workload driver is LoadRunner. The workload consists of the following transactions.

Table 6. Test Application Workload

TRANSACTION NAME	DETAILS	% USERS
1. S1assign	The loan admin picks up a task from the queue and assigns to a loan agent – new_app1_x.	8
2. S2finish	Agent new_app1_x logs on to TaskSpace, acquires and finishes the task.	40
3. S3mon	The supervisor slc_monitor logs on, acquires the task, set attributes and modifies the task.	8
4. S4assign	The admin assigns the modified task in the queue to agent new_app2_x .	4
5. S5finish	Agent new_app2_x logs on to TaskSpace, acquires and finishes the task.	40

Observed Results

The graphs below show Average and Max CPU utilizations, and memory consumptions for 50, 100, 200 and heavy user tests. The 50 and 100 user tests involved one instance of Task Space Server, the 200 user test involved 2 instances of Task Space on 2 virtual machines, and the 50 heavy user test involved 2 instances of Task Space on one virtual machine. The detailed configuration of the virtual machines is specified in the Virtual Machine Configuration section provided earlier.

The results show, especially in the TaskSpace CPU utilization chart, a significant difference between Maximum and Average CPU utilization. This suggests two observations:

1. The virtual machine configuration used in the tests adequately handles load at the various user levels and transaction loads.
2. As seen in the Heavy User (high transaction rate) scenario it is possible to max out the CPU utilization. This situation will cause intermittent response time increase but can be alleviated through horizontal scaling. For planning purposes, administrators should leave adequate head-room for workload spikes.

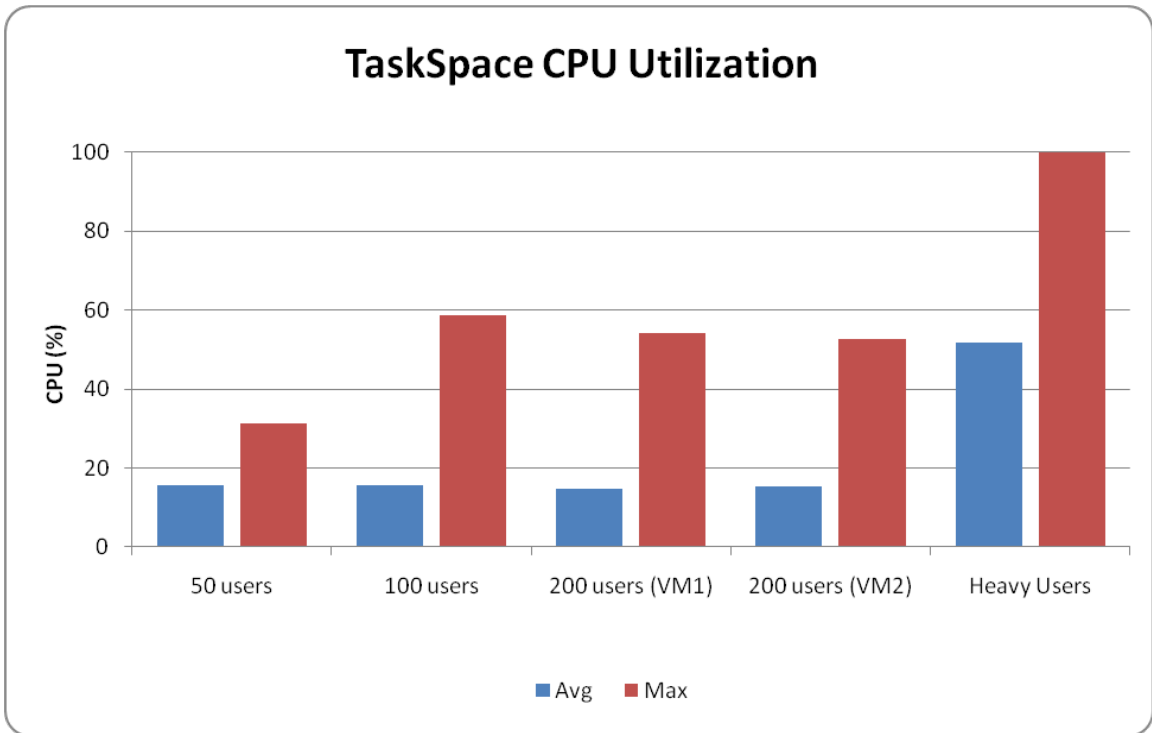


Figure 4. TaskSpace CPU Utilization

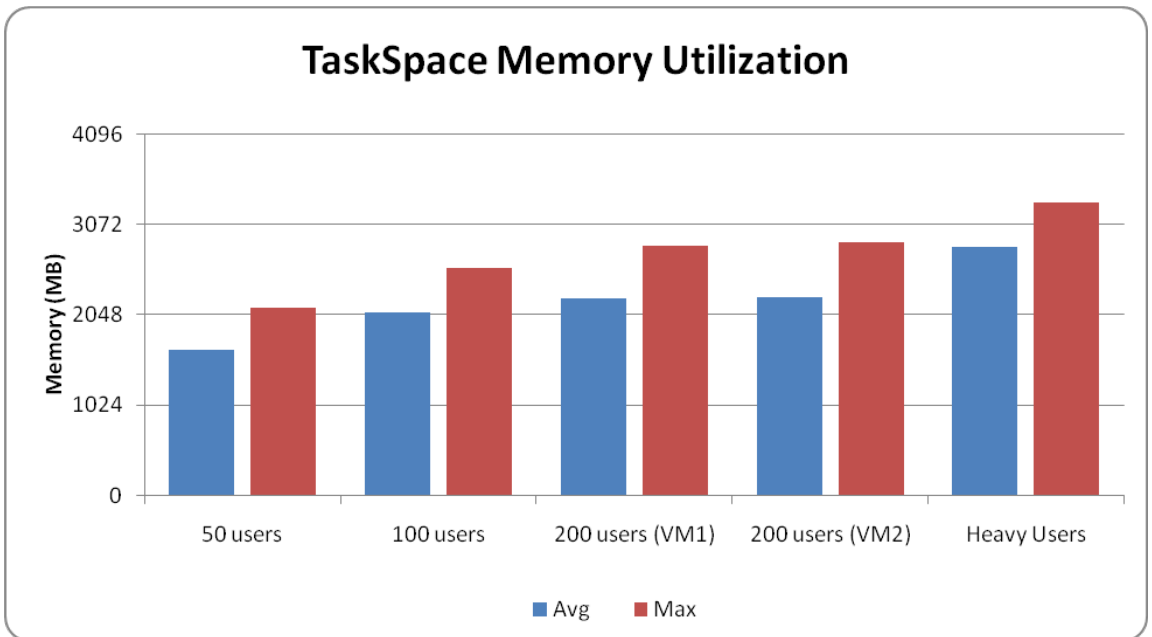


Figure 5. TaskSpace Memory Utilization

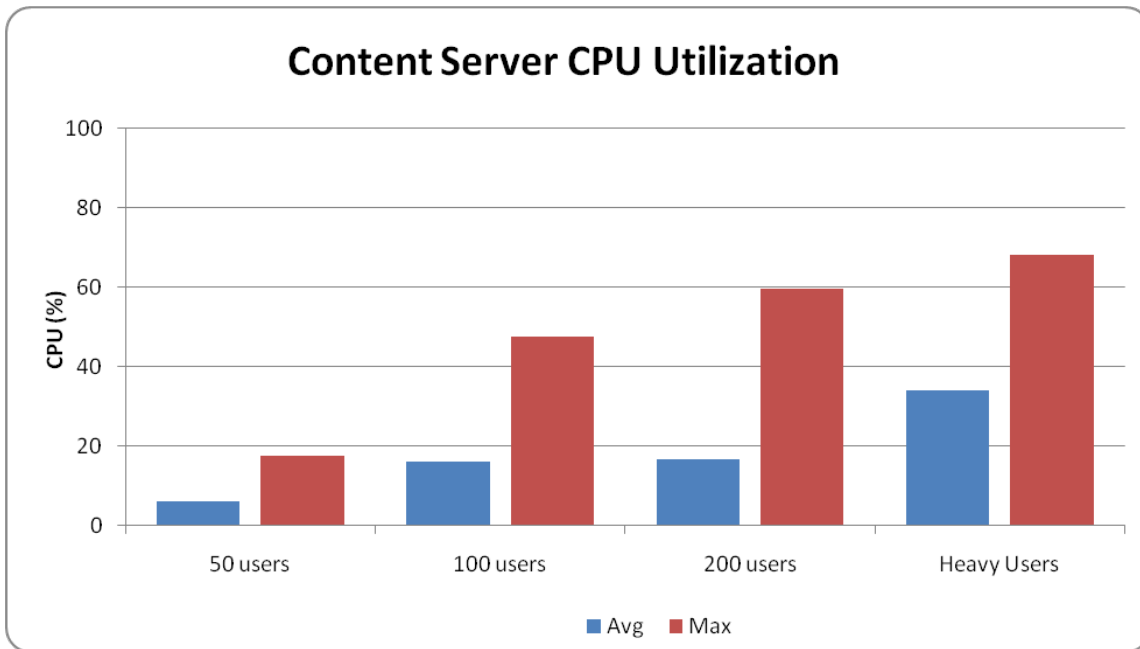


Figure 6. Content Server CPU Utilization

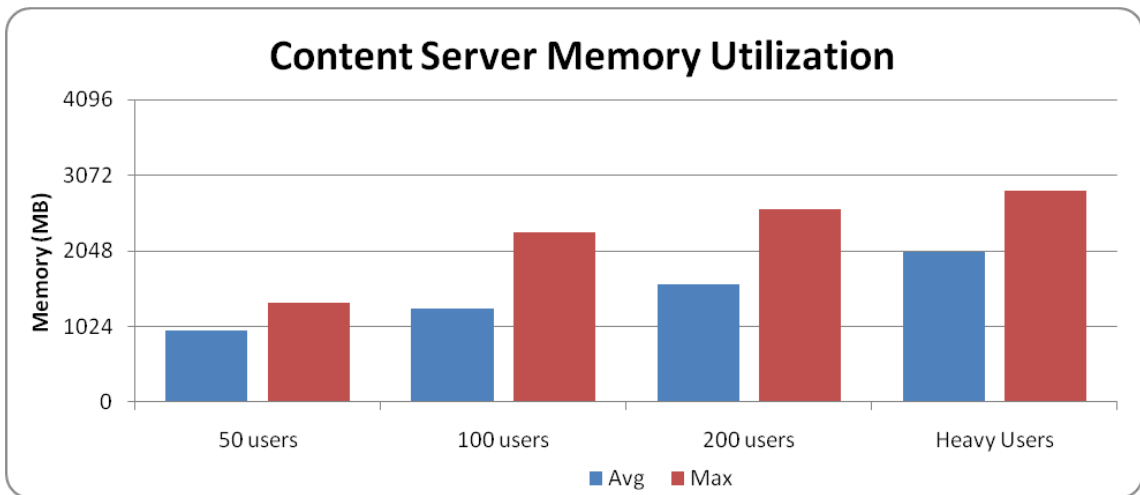


Figure 7. Content Server Memory Utilization

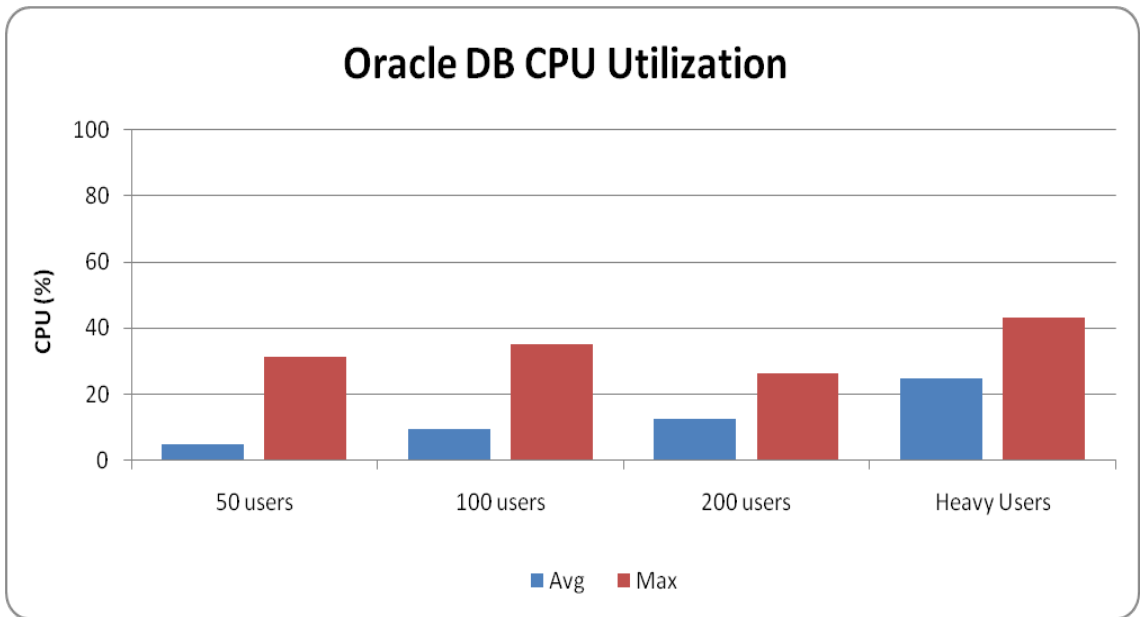


Figure 8. Oracle DB CPU Utilization

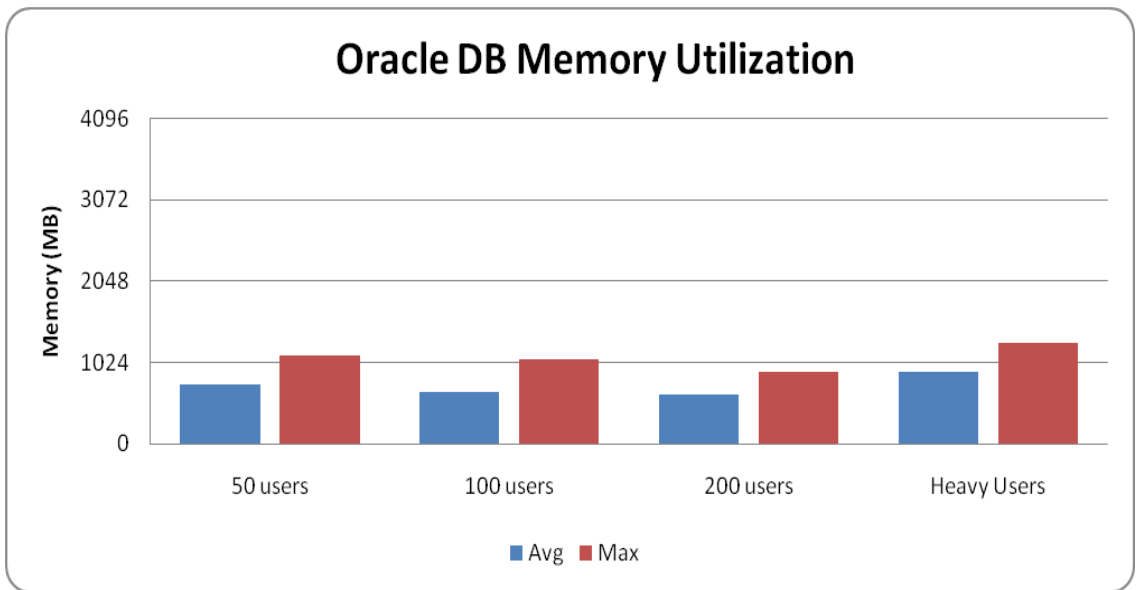


Figure 9. Oracle DB Memory Utilization

Business Continuity and High Availability Testing

High Availability

VMware HA delivers the availability needed by many applications running in virtual machines, independent of the operating system and application running in the virtual machine. VMware HA provides uniform, cost-effective failover protection against hardware and operating system failures within your virtualized IT environment. VMware HA uses the following approach to provide increased business continuity for workloads running on the vSphere platform.

- Monitors virtual machines to detect operating system and hardware failures.
- Restarts virtual machines on other physical servers in the resource pool, without manual intervention, when server failure is detected.
- Protects applications from operating system failures by automatically restarting virtual machines when an operating system failure is detected.

High Availability Test

To examine how TaskSpace and Content Server could leverage VMware HA, the test arranged TaskSpace and Content Server virtual machines on one host (ESX host-A) of a three node cluster. The physical host for ESX-A was then rebooted. VMware vSphere detected the interruption and automatically restarted the TaskSpace virtual machine on ESX host-B while it restarted the Content Server virtual machine on ESX host-C. The automatic restart completed in 180 seconds.




Figure 10. High Availability Test for TaskSpace and Content Server

High Availability Observations for TaskSpace and Content Server

Initially, after the automatic restart, approximately 50% of TaskSpace transactions began falling into an error state. This likely occurred due to the fact that the TaskSpace virtual machine booted before Content Server. As a result, the TaskSpace instance could not properly service incoming transaction requests.

However, after manually restarting the TaskSpace process in the virtual machine incoming transactions appeared to resume executing correctly. This observation highlights the opportunity to leverage a vApp construct to better orchestrate virtual machine initialization.

VMotion

VMotion enables the live migration of running virtual machines from one physical server to another, with zero downtime, continuous service availability, and complete transaction integrity. This capability makes hardware maintenance possible at any time of the day and does not require clustering or redundant servers. VMotion makes it possible to move online workloads as required, from one ESX host machine to another, in order to maintain service levels and performance. The scope of this test was to check the migration of the AppServer virtual machine.

vMotion Test

With the TaskSpace and Content Server virtual machines running in separate ESX hosts (after the High Availability test was performed), the vMotion test initiated the migration of both virtual machines back to the initial ESX host-A.




Figure 11. vMotion Test with TaskSpace and Content Server Virtual Machines

vMotion Observations for TaskSpace and Content Server

VMotion migration was completed in 110 seconds, with no failures, indicating that user activity could continue successfully during any VMotion events. No negative impact on transaction execution time was observed.

Deployment Best Practices

VMware recommends the following best practices derived from configuration and testing of the Documentum xCP application running on VMware virtual infrastructure.

- It is recommended that you use Intel-based Nehalem or AMD-based Shanghai processors which have the Hardware-Assist capabilities that help in performing memory management operations.
- While using Intel-based systems, ensure that the BIOS settings enable Hyperthreading, VT, and EPT options on all ESX hosts.
- Use Eager zeroed disks for better I/O performance. This can be done using the following VMware vmfstools command:

```
vmfstools -w abc.vmdk
```
- To decrease disk latency, it is recommended that you use Fibre Channel Adapters and configure proper RAID configuration with sufficient spindles.
- Install VMware tools on the virtual machines. The VMware Tools package provides support required for shared folders and for drag and drop operations. Other tools in the package support synchronization of time in the guest operating system with time on the host, automatic grabbing and releasing of the mouse cursor, copy and paste operations between guest and host, and improved mouse performance in some guest operating systems.
- Consider using server-class network interface cards (NICs) for the best performance and configure paravirtualized vmxnet3 adapters for better network throughput.
- Ensure that the Service Console / VMkernel and virtual machines are configured on separate NICs.
- Disconnect or disable unused or unnecessary physical hardware devices, such as:
 - COM ports
 - LPT ports
 - USB controllers
 - Floppy drives
 - Optical drives (that is, CD or DVD drives)

Disconnecting or disabling these devices will help free up interrupt resources, because traditionally some devices such as USB controllers operate using a polling scheme that consumes extra CPU resources. Lastly, some PCI devices reserve blocks of memory, making that memory unavailable to ESX.

Scaling Up to Larger Configurations

While the testing described in this document focused on servicing up to 200 concurrent users, test results may be extrapolated to system architectures and deployments that support higher user counts. The configurations tested here are not CPU-bounded. For example, in the 200-user test, content server CPU utilization peaked at around 60%, while the two TaskSpace virtual machines consumed a total of 30% CPU, on average, and each peaked at around 55%. As a matter of fact, TaskSpace capacity is limited by Java heap size — a 1.5GB heap is allocated for each TaskSpace instance on the two virtual machines for the 200-user test.

For the heavy user test, TaskSpace is configured differently, with two instances of TaskSpace running on a single virtual machine. In this case, we scale out TaskSpace within the virtual machine. Since the transaction rate is approximately twice the 200-user test, average CPU utilization is around 52%, but utilization also spikes to 100%. This indicates that a single virtual machine can host 200 users with two TaskSpace instances on a 32-bit OS, or one instance with a large heap size (3GB or above) on a 64-bit OS.

For cases beyond 200 users, a “cookie cutter” approach could be used to define suitable deployment environments, where components of the 200-user system architecture are replicated N times to support (N x 200) users. Naturally, the scaling factor is likely to vary for xCP applications with different levels of complexity and workflows, and care should always be taken that the system parameters are adjusted correctly. For more information please refer to the *Documentum xCP Best Practices Guide*.

The following two figures show the standard system blueprint used for 50 users and how that blueprint can be scaled to meet the needs of 1000 concurrent users.




Figure 12. Basic TaskSpace Blueprint

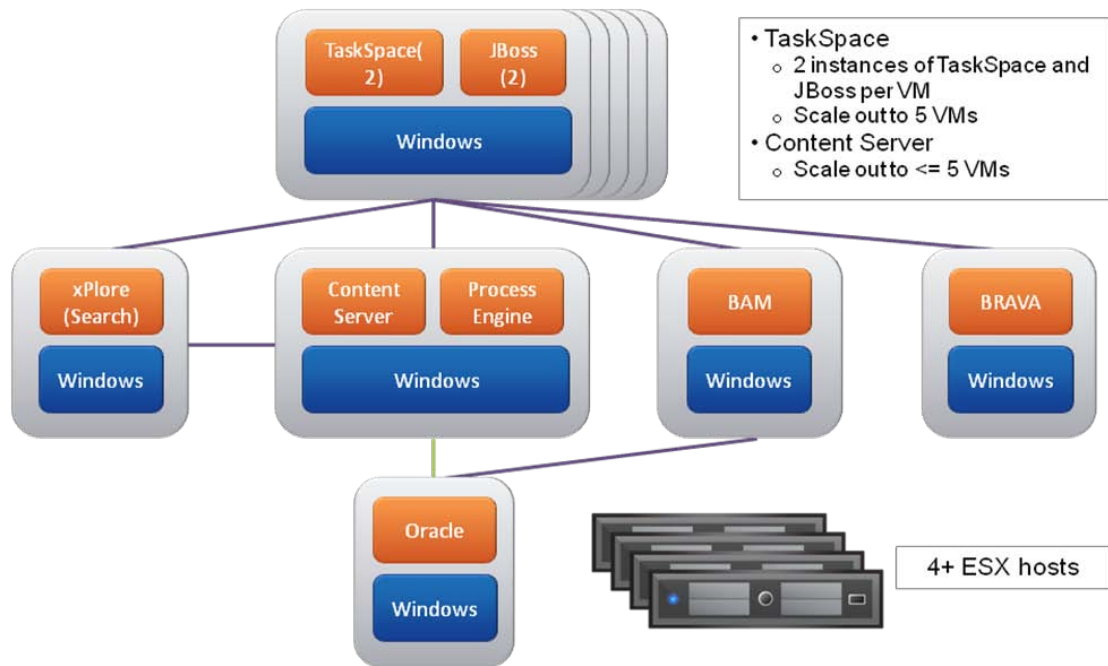


Figure 13. TaskSpace Blueprint at Scale

Leveraging the nicely linear scalability of TaskSpace and Content Server on VMware vSphere to meet projected load at 1000 concurrent users, we would instantiate two TaskSpace processes in each virtual machine and scale that virtual tier horizontally with five virtual machines. We would also scale the Content Server tier, expecting to use five or fewer virtual machines. To accommodate the growing number of virtual machines, we would add at least one ESX host and monitor performance to see if further hosts are needed.

Summary and Conclusions

Overall, testing results show that running Documentum xCP on VMware virtual infrastructure performs well within expectations based on various user and transaction loads. Documentum xCP meets additional load with near 1:1 linear horizontal scalability, which enables very predictable growth and capacity planning.

With respect to business continuity, Documentum xCP leverages vMotion very well for “planned downtime” scenarios. Administrators can migrate live Content Server and TaskSpace virtual machines to allow for ESX host maintenance without impacting end-user uptime or performance. Having a standard procedure in place for initialization orchestration allows Documentum xCP to take full advantage of VMware HA for “unplanned downtime” scenarios.

EMC’s Documentum xCP platform, including the xPlore Search capability, exhibits performance, scalability, and business continuity behavior worthy of VMware Ready™ Certification. Customers and partners should investigate the benefits of virtualization using the vSphere platform and feel assured that EMC Documentum will meet their needs when running in a VMware virtual infrastructure environment.

Additional Resources

Customers can find more information about VMware and EMC Documentum xCP products via the links listed below.

VMware Resources

- VMware official website:
<http://www.vmware.com/>
- VMware Infrastructure 3 and vSphere 4 product Web site:
http://www.vmware.com/products/data_center.html
- VMware download Web site:
<https://www.vmware.com/download/>
- VMware support Web site:
<http://www.vmware.com/vmtrn/>
- VMware Performance Tuning Paper:
http://www.vmware.com/pdf/vi_performance_tuning.pdf
- System Compatibility Guide for a complete list of compatible hardware:
http://www.vmware.com/pdf/vi35_systems_guide.pdf
- Storage/SAN Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/pdf/vi35_san_guide.pdf
- I/O Compatibility Guide for a complete list of compatible networking devices:
http://www.vmware.com/pdf/vi35_io_guide.pdf

EMC Documentum xCP Resources

- EMC Documentum xCP Web site:
<http://www.emc.com/xCP>
- EMC Documentum xCP Product Family Web site:
<http://www.emc.com/products/family/documentum-xcp-family.htm>
- EMC Documentum Developer Community:
<http://developer.emc.com/xCP>
- EMC Documentum xCP Best Practices Guide:
<https://community.emc.com/docs/DOC-4096>
- EMC Documentum xCP - A Detailed Review Whitepaper:
<http://www.emc.com/collateral/software/white-papers/h4723-documentum-xcp-wp.pdf>
- EMC Press Release:
<http://www.emc.com/about/news/press/2010/20101110-01.htm>

