

Virtualizing Mission-Critical Applications

Following the successful proof of concept, Intel IT plans to start virtualizing the first of many mission-critical applications in 2011 as part of our private cloud environment.

Nickie Lee
EAI Architect
Intel IT

Uttam Shetty
Private Cloud Engineering
Intel Architecture Group

Raghuram Yeluri
Private Cloud Architect
Intel Architecture Group

Executive Overview

To create the infrastructure for a private enterprise cloud, Intel IT has set a goal of virtualizing up to 75 percent of our office and enterprise computing environments. To achieve this goal, Intel IT will need to virtualize mission-critical applications, which is challenging due to rigorous performance, availability, and other requirements.

To determine the feasibility of virtualizing mission-critical applications and identify optimum server consolidation ratios, Intel IT and Intel Architecture Group conducted proof-of-concept (PoC) testing. The testing focused on an enterprise application integration (EAI) platform used to exchange confidential, sensitive, and essential business information between Intel and its customers and suppliers. We compared performance, availability, and deployment time with physical and virtualized configurations. We found that the virtualized configuration:

- Delivered 1.4x the maximum throughput of the physical configuration and 5x the average throughput observed in the production environment, demonstrating the ability to meet performance service-level agreements (SLAs) and accommodate anticipated growth. We also improved virtualization performance

by up to 40 percent by enabling Intel® Hyper-Threading Technology.

- Achieved a 5:1 server consolidation ratio in the tests.
- Greatly reduced unplanned downtime due to server failure; downtime was five minutes in the virtualized environment compared to 30 minutes in the physical environment.
- Could more quickly accommodate increases in workload by deploying new EAI components in four hours, compared with two weeks in the physical environment.

Following the successful PoC, Intel IT plans to start virtualizing this EAI platform, the first of many mission-critical applications, in 2011.

The PoC opens the door to virtualization of other mission-critical applications at Intel and represents a significant step toward the goal of virtualizing 75 percent of the Intel IT office and enterprise environments.

Contents

Executive Overview.....	1
Business Challenge	2
Proof of Concept.....	2
PoC Objectives.....	3
PoC Architecture	3
Results	4
Next Steps	6
Conclusion	6

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BUSINESS CHALLENGE

Intel IT is implementing a private cloud based on flexible, energy-efficient virtualized infrastructure to realize cloud computing benefits such as increased agility and reduced cost. As a first step to building our private cloud, we are rapidly virtualizing our existing office and enterprise computing environments, with the goal of virtualizing up to 75 percent of applications.

We have focused initial efforts on applications that lend themselves most readily to virtualization, yielding the greatest benefit with the lowest risk. Examples include departmental and line-of-business applications. We are achieving consolidation ratios of up to 20:1 when virtualizing older servers running these applications and consolidating them onto servers based on the latest Intel® Xeon® processors.

In order to achieve the ultimate goal of virtualizing up to 75 percent of our environment, we will eventually have to virtualize mission-critical applications as well. This is perceived as much more challenging, with significant risks. Mission-critical applications must meet very high performance and availability requirements, defined in service-level agreements (SLAs) with Intel's business groups.

Performance is a key concern. Because virtualization involves moving applications from dedicated physical servers to a shared environment, there is a risk of performance degradation due to the need to share storage and network resources.

Intel IT, with Intel Architecture Group, conducted proof-of-concept (PoC) testing to assess the feasibility of virtualizing mission-critical applications within Intel and to determine whether the potential

benefits outweighed the risks. Another key objective was to establish optimum server consolidation ratios. Our testing focused on areas such as performance, availability, speed of application deployment, and total cost of ownership (TCO).

PROOF OF CONCEPT

For our PoC, we focused on a mission-critical enterprise application integration (EAI) platform. This enables and automates the exchange of essential business documents, such as purchase orders, between Intel's internal business applications and Intel's customers and suppliers. It accepts information from these trading partners, translates the data into formats Intel uses internally, and then passes it to our enterprise resource planning (ERP) and other internal applications. The platform also translates information passing in the reverse direction—from internal applications to Intel customers and suppliers. A high level-overview is shown in Figure 1.

Because the platform handles information that directly affects Intel's revenue, its performance and availability are critical. Each business document generates one or more messages that must be processed by the EAI platform. The platform currently handles about 1 million messages per day, and this number is expected to increase over time to 3 million messages per day. The number of users is expected to increase from about 600 to approximately 2,000.

This mission-critical platform has a 99.99 percent uptime requirement. It processes many different types of information from multiple trading partners; the SLAs for

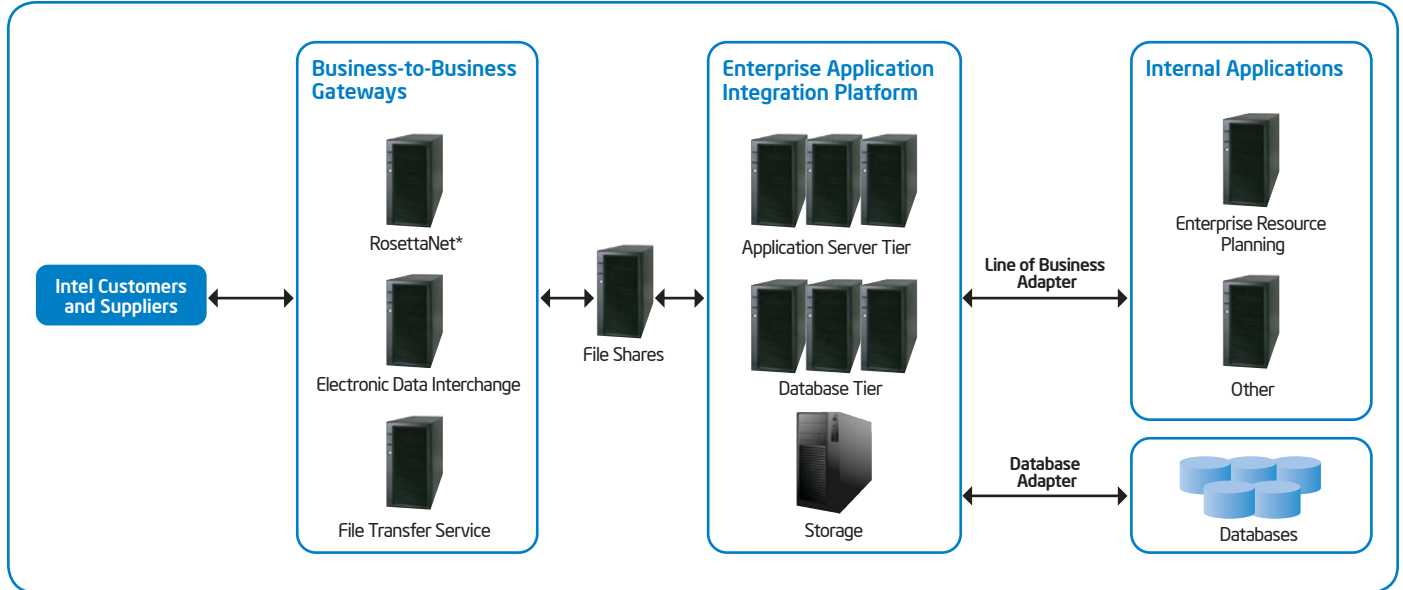


Figure 1. Role of the mission-critical enterprise application integration platform within the Intel IT environment.

some of these transactions stipulate that messages must be processed and transmitted within five minutes of receipt.

PoC Objectives

A key goal of the PoC was to determine whether the benefits of virtualizing a mission-critical platform outweigh the risks. We focused on several specific areas.

FEASIBILITY OF VIRTUALIZING MISSION-CRITICAL APPLICATIONS

We needed to determine whether a virtualized implementation could:

- Achieve performance equal to or better than the physical production environment.
- Accommodate the expected growth in message volume.
- Reduce unplanned and planned downtime.

RETURN ON INVESTMENT AND AGILITY

We assessed additional potential benefits such as faster application deployment and reduced cost due to server consolidation.

CONSOLIDATION RATIOS

We also designed our testing to determine server consolidation ratios that optimize performance and resource utilization. If our testing was successful, the design could be used as a template for virtualizing other EAI instances within Intel. EAI platforms are used for many purposes at Intel, running on a total of about 100 physical servers.

PoC Architecture

The EAI platform includes an integration engine consisting of three application server tiers (receiving, transmitting and orchestration). The platform also includes a database tier.

- **Receiving tier.** Processes incoming messages and stores them in one or more message boxes located on a database server.
- **Transmitting tier.** Retrieves messages from a message box and routes them to destination systems.
- **Orchestration tier.** Defines and applies actions to each message in order to implement specific business processes.

The implementation that we tested in the PoC is complex: It handles more than 100 different document types from multiple trading partners.

- **Database tier.** In addition to the application tiers above, the platform includes databases that store the message boxes and track progress.

Each of these tiers is deployed on one or more servers. Within each tier, throughput can be increased using a scale-out approach—creating new instances of the EAI component and deploying them on additional servers.

TEST ENVIRONMENT

For the PoC, we created a physical test configuration that matched our production configuration. As shown in Figure 2, this consisted of the application and database tiers deployed across 10 physical servers, each based on Intel® Xeon® processor 5400 series.

We compared this physical setup with a virtualized configuration on a dedicated network representing the Intel private cloud environment, also shown in Figure 2.

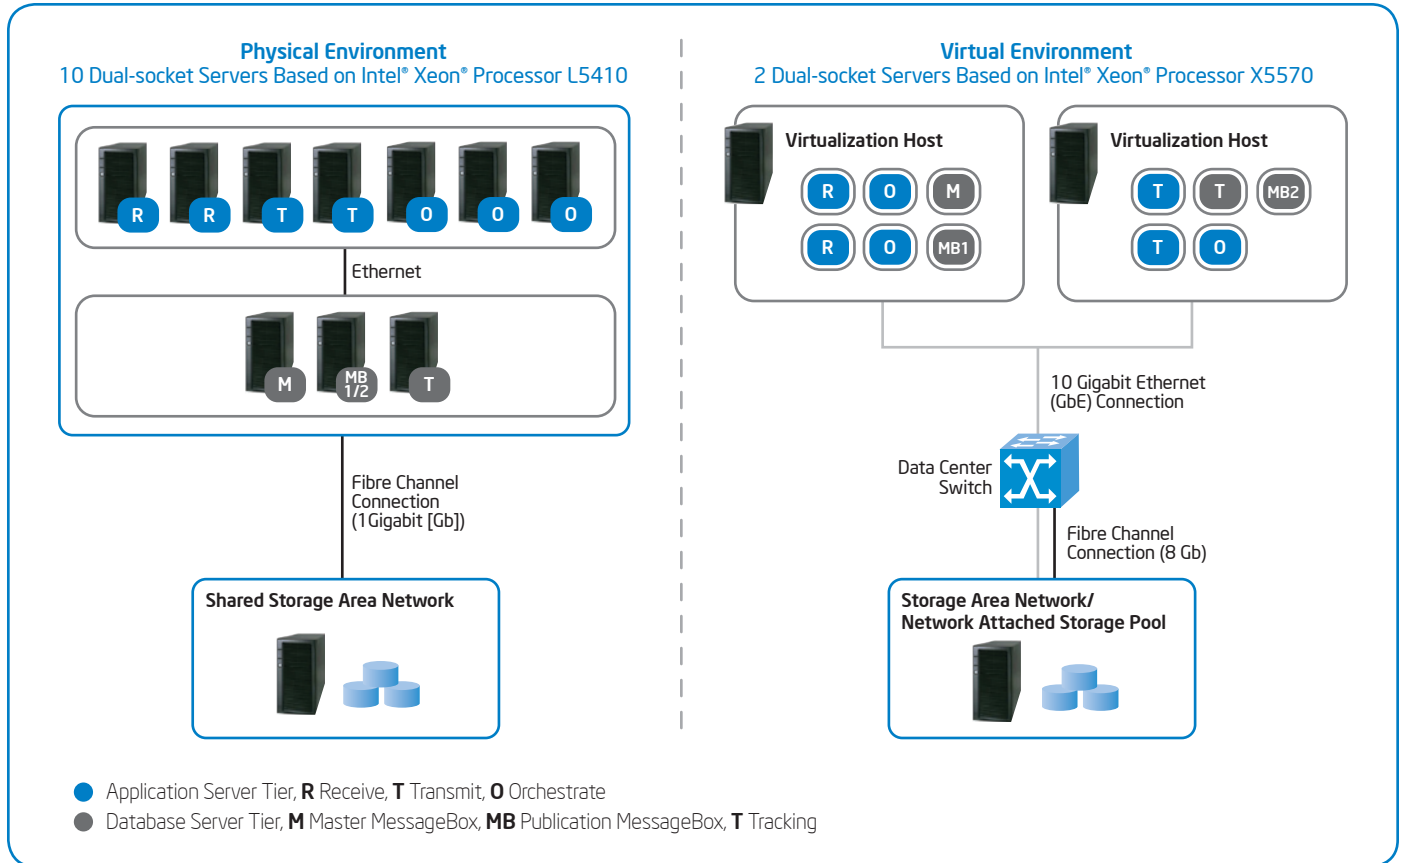


Figure 2. Proof of concept architecture.

To create this configuration, we virtualized each of the application and database servers and consolidated them onto two physical virtualization host servers. This represented a 5:1 consolidation ratio.

Each of the virtualization hosts was a dual-socket server based on Intel® Xeon® processor 5500 series, which includes features that increase virtualization performance such as Intel® Virtualization Technology (Intel® VT) and Intel® Hyper-Threading Technology (Intel® HT Technology).

We determined the resources that should be allocated to each virtualized server by analyzing utilization of the corresponding physical server in the current production environment. Detailed specifications are shown in Table 1.

TEST PROCEDURE

We created test scripts to generate representative transaction workloads. We compared peak and average throughput in the physical and virtualized test environments.

We also simulated planned and unplanned downtime scenarios, and compared the time required to deploy new server instances to increase throughput and accommodate demand growth.

Preliminary testing indicated that enabling Intel HT Technology in the virtualization host servers based on Intel Xeon processor 5500 series increased throughput by up to 40 percent. We therefore enabled Intel HT Technology on these servers throughout the PoC.

Results

In our tests, the virtualized configuration provided greater throughput than the physical configuration, demonstrating the capacity to handle anticipated production workload growth. Tests also indicated that the virtualized configuration could reduce downtime and accelerate deployment. Results are summarized in Table 2.

PERFORMANCE AND UTILIZATION

Maximum throughput of the virtualized configuration was 436,000 messages per hour, 1.4x the maximum throughput of 259,000 messages obtained using the physical setup and about 5x the average throughput (86,000) in the current production environment.

Table 1. Proof of Concept System Specifications

		Physical Configuration		Virtual Configuration
		Specification	Utilization	Resources Allocated
Application Server Tier (Specifications for each server)	CPU	<ul style="list-style-type: none"> Dual-socket servers with 2x Intel® Xeon® processor L5410 (2.33 GHz, quad-core) 	<ul style="list-style-type: none"> 15 to 20% of 8 cores = 2 cores 	<ul style="list-style-type: none"> 2 virtual CPUs on a dual-socket server with 2x Intel® Xeon® processor X5570 (2.93 GHz, quad-core)
	Memory	<ul style="list-style-type: none"> 8 GB 	<ul style="list-style-type: none"> 32-bit OS; maximum use 5.5 GB 	<ul style="list-style-type: none"> 72 GB; 6 GB per virtual machine (VM)
	Network	<ul style="list-style-type: none"> 3x 100 Mb network interface cards (NICs) 	<ul style="list-style-type: none"> 10 Mb – production 50 Mb – backup 5 Mb – remote management 	<ul style="list-style-type: none"> 1x 1-Gb NIC 1x 10-Gb NIC
	Local Storage	<ul style="list-style-type: none"> 40 GB 	<ul style="list-style-type: none"> 40 GB 	<ul style="list-style-type: none"> None
Database Tier	CPU	<ul style="list-style-type: none"> Dual-socket servers with 2x Intel Xeon processor L5410 (2.33 GHz, quad-core) 	<ul style="list-style-type: none"> 52% of 8 cores = 4 cores 	<ul style="list-style-type: none"> 2 virtual CPUs on a dual-socket server with 2x Intel Xeon processor X5570 (2.93 GHz, quad-core)
	Memory	<ul style="list-style-type: none"> 8 GB 	<ul style="list-style-type: none"> 32-bit OS; maximum use 5.5 GB 	<ul style="list-style-type: none"> 72 GB; 6 GB per VM
	Network	<ul style="list-style-type: none"> 100 Mb 	<ul style="list-style-type: none"> 20 Mb – production NIC 20 Mb – backup 20 Mb – remote management 	<ul style="list-style-type: none"> 1x 1-Gb NIC 1x 10-Gb NIC
	Storage	<ul style="list-style-type: none"> 2x 2 Gb host bus adapter (HBA) (for SAN access) Local storage: 20 GB 	<ul style="list-style-type: none"> Storage IOPS: 150 to 200 on average; peak:1000 	<ul style="list-style-type: none"> 2x 8 Gb HBA (for SAN access)

Based on this result, the virtualized configuration could accommodate the anticipated growth in the production workload to 3 million messages per day, or 375,000 messages per hour.

The servers were utilized more efficiently in the virtualized environment. CPU utilization on the virtualization host servers based on Intel Xeon processor 5500 series averaged 35 to 42 percent, reaching a maximum of 56 percent. This was comparable with utilization of 15 to 20 percent for application servers and 52 percent for database servers in the physical configuration.

AVAILABILITY

Testing indicated that virtualization could significantly reduce unplanned downtime. It could also eliminate the need for planned downtime when server maintenance is required.

PLANNED DOWNTIME

In the physical environment, server maintenance requires application downtime, which must be scheduled in advance with business groups. In a virtualized environment, the hosts running the EAI platform are part of a virtualization cluster, and live migration can be used to move the EAI platform off one host and onto another host within the cluster. This provides the ability to perform maintenance on the original host without the need for application downtime.

During the PoC, we validated this approach by seamlessly moving four virtual machines (VMs) off one of the host servers and onto another server. We verified that transactions continued to run during the migration, without interruption to service.

Unplanned Downtime

In the physical environment, a database or application server failure impacts performance as workloads are redirected to the remaining servers in an application or database cluster. In addition, a failure can require rebooting physical servers, which can take up to 30 minutes. In the virtualized environment, in the event of a crash, we can restore VMs to another server, reducing the unplanned downtime required. For the PoC, we simulated a crash by powering off VMs using the virtualization software management console. Successfully restoring all the VMs took five minutes—one-sixth the time required for a physical server reboot.

AGILITY

In both the physical and virtualized environments, it is possible to accommodate

Table 2. Summary Results

	Physical Servers	Virtual Servers	Improvement
Number of Servers Required	10	2	80% reduction
Maximum Throughput (Messages per hour)	259,000	436,000	1.4x increase
Average Utilization	15 to 20% for application servers; 52% for database servers	35 to 42%	Approximately 2x improvement for application servers
Approximate Time to Restart Application or Database Server Following Failure	30 minutes	5 minutes	83% reduction
Time to Deploy New Platform Component	80 hours	4 hours	95% reduction

demand growth by scaling out—adding more servers within a tier. In the virtualized environment, this can be completed much more quickly, in about four hours compared with two weeks in the physical environment.

Intel IT application requirements are complex; the platform includes about 100 different interfaces to handle information from different trading partners. In a physical environment, deploying all these interfaces on a new server is time-consuming and labor-intensive; because the interfaces are provided by a variety of different developers, we cannot automate the process with a single script. As a result, it can take two weeks—80 work hours—deploy a new physical server and install all these interfaces.

In the virtual environment, it is possible to simply clone a new VM from an existing VM, which guarantees that all components are installed in the new VM. In our tests, this required approximately four hours.

NEXT STEPS

Due to the success of our PoC, we plan to virtualize the mission-critical production EAI application in 2011, once the necessary private cloud infrastructure is in place and we have conducted further performance testing in a shared network and storage environment.

We are analyzing the potential for further optimization using Intel® architecture features such as Intel® Turbo Boost Technology, included in the latest Intel Xeon processors, and Intel® Virtualization Technology Extended Page Tables, a feature of Intel VT that can increase virtualization performance with memory-intensive application workloads.

CONCLUSION

Our PoC demonstrated that virtualizing business-critical applications is not only feasible but offers considerable benefits that outweigh the risks.

In our tests, the virtualized configuration delivered greater throughput than the current physical production configuration, demonstrating that the virtualized platform can meet existing SLAs and provide headroom for growth. Virtualization performance improved up to 40 percent using Intel HT Technology.

The virtualized configuration delivered this performance while achieving a 5:1 consolidation ratio compared to the current physical production configuration, demonstrating the potential to reduce server TCO. We also observed significant improvements in availability and the speed of application deployment.

The success of our PoC is a significant step toward the ultimate goal of virtualizing up to 75 percent of the Intel IT environment. We plan to move ahead with virtualizing this EAI application during 2011, and we are also preparing to virtualize other EAI servers. The successful tests also open the door to virtualization of other mission-critical applications at Intel, including externally facing applications.

For more information on Intel IT best practices, visit www.intel.com/it.

CONTRIBUTORS

Chris Black

Private Cloud Infrastructure Engineer, Intel IT

Ajay Chandramouly

Industry Engagement Manager, Intel IT

Vishwa Hassan

IT Research and Technology Development, Intel IT

Ashish Khare

EAI Engineering, Intel IT

Bharat Mohla

EAI Engineering Manager, Intel IT

Sanjay Rungta

Principal Engineer and Network Architect, Intel IT

Terry Yoshi

Storage Architect, Intel IT

ACRONYMS

EAI	enterprise application integration
ERP	enterprise resource planning
Gb	gigabit
GbE	gigabit Ethernet
HBA	host bus adapter
Intel® HT Technology	Intel® Hyper-Threading Technology
Intel® VT	Intel® Virtualization Technology
NIC	network interface card
PoC	proof of concept
SLA	service-level agreement
TCO	total cost of ownership
VM	virtual machine

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/performance/resources/benchmark_limitations.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE

ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2011 Intel Corporation. All rights reserved.

Printed in USA
0111/ABC/KC/PDF

 Please Recycle
324168-001US

