

IT@Intel Brief

Intel Information Technology

Computer Manufacturing
Server Virtualization

March 2007

Virtualizing the Data Center

Intel IT is implementing a major data center virtualization (DCV) initiative that enables Intel design engineers to transparently use compute server resources worldwide, increasing server utilization and reducing costs.

We have already achieved an 11 percent increase in server utilization and saved USD 77 million through cost avoidance. Our goal is to increase design computing server utilization to 80 percent or more.

DCV is a long-term initiative that includes re-engineering business processes, implementing new technologies, and adapting user workloads. Virtualization is a key part of our overall data center optimization strategy, which also includes consolidating data centers and accelerating server refresh cycles.

Profile: Virtual Data Center

- Server utilization increased from 55 to 66 percent
- USD 77 million saved to date
- Goal is to increase server utilization to 80 percent

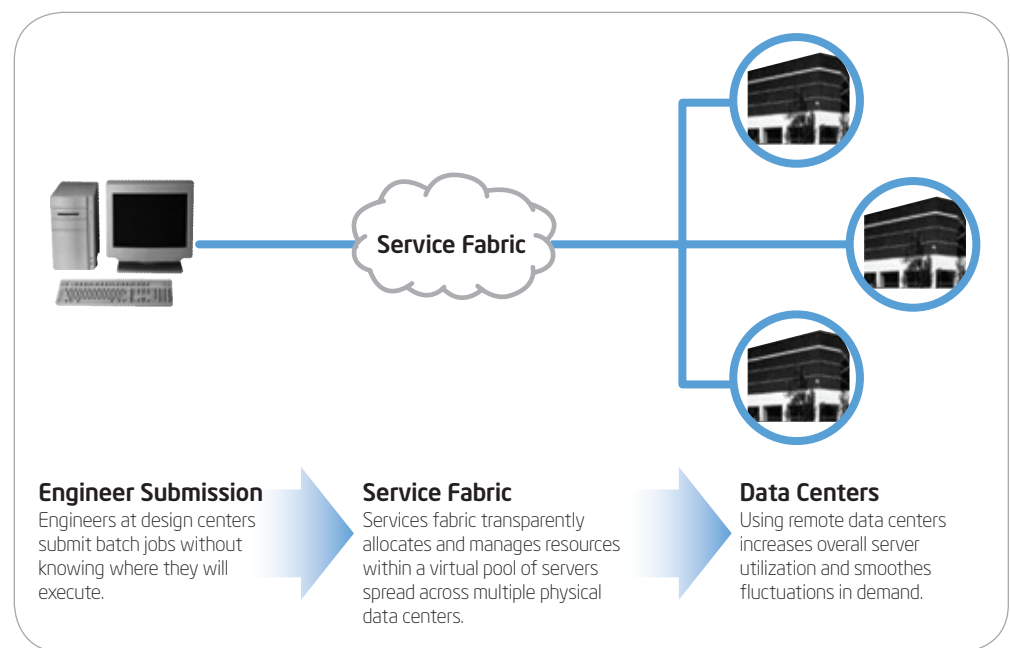


Figure 1. Data center virtualization.

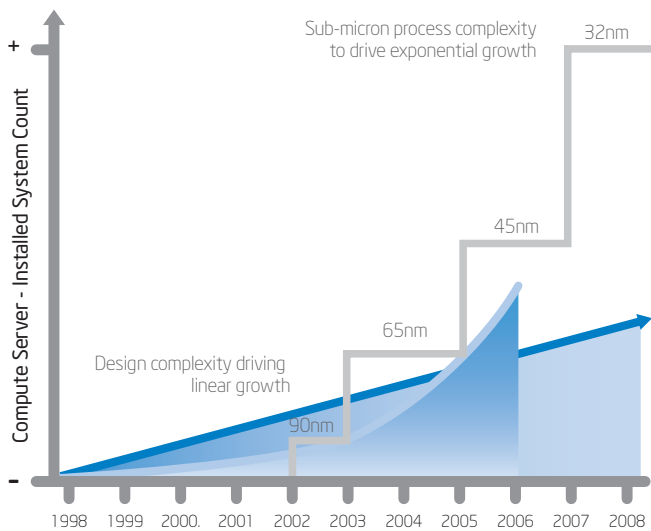


Figure 2. Increasing design and process complexity requires more compute power.

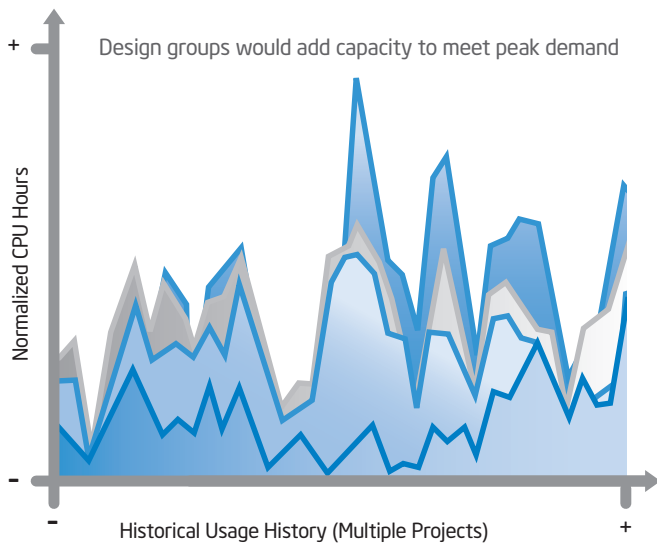


Figure 3. Local fluctuations in server demand. We see similar fluctuations in use over multiple projects.

Background

Designing semiconductors requires ever-increasing amounts of computing power, as chip designs and manufacturing processes become more complex. Most of Intel's servers support design computing and the number of these servers has increased rapidly, surging fifty-fold to more than 60,000 between 1996 and 2005, as shown in Figure 2.

Traditionally, we have installed many of these servers in local design data centers co-located with design teams. Intel has more than 70 of these data centers worldwide.

To obtain the growing computing power needed for their projects, design groups typically added enough local capacity to meet peak demand, as shown in Figure 3. This meant that in between the demand peaks, many of these local servers were underutilized. Worldwide, this added up to significant underused capacity.

We are addressing this issue by virtualizing data center resources. By creating processes and infrastructure that let design teams use idle servers at remote data centers, we can apply Intel's global computing resources to individual projects. This increases overall server utilization and smoothes fluctuations in demand, reducing the need to add capacity and cutting costs. It potentially also accelerates design steps and improves the quality of Intel® products.

Virtualization Approach

Our goal is to make using remote resources highly automated and transparent, using an approach based on open service-oriented architecture (SOA) standards and grid computing concepts. Design teams should be able to use computing resources anywhere, without being concerned about their physical location. While not all design activities can benefit from DCV, those that do account for about 50 percent of all compute cycles and 70 percent of the batch demand. These tasks typically consist of large compute-intensive batch jobs that execute on many servers simultaneously. With DCV, an engineer at a local site could submit a batch job that would transparently execute on a virtual pool of servers potentially spread across multiple physical data centers.

Initially, however, it required considerable effort and time to set up and run design environments remotely. Some project environments include millions of lines of scripting that we needed to adapt to run the environments at multiple remote sites. We also lacked tools to enable DCV operation and management. As a result, it required 35 days on average to set up a project for remote execution and the effort involved limited our ability to share resources. We realized that we needed to overcome these barriers by developing or adapting operations, business processes, technologies, and user workloads.

Together with design teams, we identified five key business processes that we needed to re-engineer in order to support our goals:

- **Capacity planning** was difficult and not very accurate, requiring better data quality and metering.
- **Resource allocation** issues centered on ownership, prioritization, and license management, with a need for reservation tools and processes.
- **Workload flow and setup** needed standardization in areas including account creation and management, disk provisioning, data replication, and OS provisioning.
- **Scripting and data synchronization** needed improvement in order to run large jobs, in areas including job monitoring, batch tools, and management of virtual pools.
- **Real-time monitoring and day-to-day operations** needed better support for remote sites and standard processes for "tear-down" of virtualized environments after completion.

During 2006, we re-architected these processes to support a virtualized environment and began using the new standardized processes across participating sites. One major change was the establishment of a worldwide council responsible for allocating available resources at these sites among the projects that need them.

Common Services

We used our re-engineered processes as a framework for aligning development of common services that support virtualization. Today we are developing most of the services ourselves, but as DCV becomes more widely adopted, we hope to be able to buy more off-the-shelf products. Current and planned services, shown in Figure 4, include:

- **Batch scheduler.** Our batch scheduler originally assigned jobs to physical servers within a data center. We have enhanced it to handle virtual server pools spread across multiple data centers and we are now adding more manageability features.
- **Indicator system.** This system applies data mining and analysis to real-time and historical data, letting us identify utilization trends, workload growth, and success rates.
- **Bare-metal provisioning.** We have developed a tool that automatically provisions OSs to as many as 100 servers in 30 minutes. We aim to enhance this to enable more rapid and flexible provisioning for large jobs, potentially provisioning as many as 1,000 servers within 15 minutes.
- **Authentication and identity management.** Our infrastructure grew from many separate UNIX*-based Network Information Service (NIS) domains. This hinders resource sharing between sites and we are working on a unified framework.

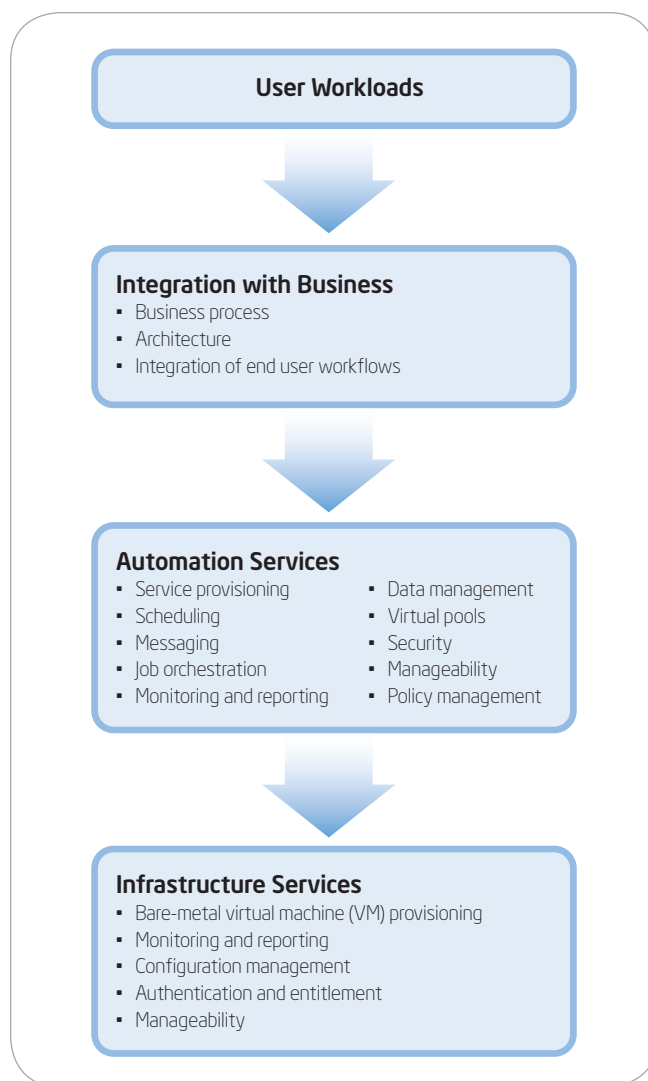


Figure 4. Common services required to support data center virtualization (DCV).

DCV Case Study

In a recent proof of concept, our Israel design team ran compute tasks on servers located in data centers in several locations worldwide. By relying on 6,000 unused and available servers at remote locations, the team shortened the design schedule by two weeks and avoided expensive server purchases.

Project Status

During 2006, we increased our batch computing utilization from 55 to 66 percent, yielding an estimated return of USD 77 million. Our worldwide allocation council currently allocates virtualized resources on about 18,000 of our more than 60,000 servers. During 2007, we aim to deliver more of our planned services and reduce the time needed to set up a remote project environment. We are aiming to reach utilization of 80 percent or more.

DCV underpins our broader data center optimization strategy. Eliminating site dependencies enables us to further reduce costs by consolidating data centers. Our strategy also includes proactively refreshing older servers. We estimate that this overall strategy will provide an estimated return of USD 428 million net present value (NPV) by 2010.

Acronyms

DCV	data center virtualization
NIS	Network Information Service
NPV	net present value
SOA	service-oriented architecture
VM	virtual machine

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, 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, Intel. Leap ahead. and Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2007, Intel Corporation. All rights reserved.

Printed in USA
03077JLM/RDA/PDF

 Please Recycle
ITAI Number: 07-01 07b

