White Paper
Intel Information Technology
Computer Manufacturing
Server Virtualization

IT Agility through Automated, Policy-based Virtual Infrastructure

Intel's Customer Solutions Group developed a method for delivering virtualized compute power throughout the enterprise, with a utility-style framework that allocates resources on demand, based on IT policies. Following a successful proof of concept (PoC), Intel IT is beginning to implement this approach in our production environment. We expect that this technology will increase agility and hardware utilization, decrease costs, and align business and IT goals.

Parviz Peiravi, Intel Corporation

October 2006

IT@Intel

Executive Summary

Intel's Customer Services Group developed a method for delivering virtualized compute power throughout the enterprise, allocating those resources on demand based on IT policies. Following a successful proof of concept (PoC) implementation of the virtualized compute infrastructure (VCI), Intel IT has begun building a similar infrastructure for our production environment.

With a fully implemented VCI, we would be able to automatically redirect compute resources efficiently to any business unit, using a utility-like model.

VCI is a set of grid, virtualization, and service-oriented technologies that allows an enterprise to share computational resources. With a fully implemented VCI, we would be able to automatically redirect compute resources efficiently to any business unit, using a utility-like model.

By implementing VCI, we expect:

- Enhanced IT agility
- Reduced costs due to resource sharing across business units
- Improved utilization of existing grid, virtualization, and provisioning technologies
- Alignment between the goals of senior management, IT, and business users

Our virtualized infrastructure is the foundation for providing IT as a service. We plan to implement chargeback, forecasting, and modeling systems based on these technologies, enabling us to support corporate goals by accurately tracking resource use by business groups.

Contents

Background

Enterprises such as Intel increasingly recognize agility as a key competitive advantage and look to their IT organizations to achieve it. If we can rapidly repurpose IT systems to meet new computing challenges as they arise, we can achieve greater agility by addressing business needs as quickly as possible. We will also obtain maximum value from our infrastructure investments.

Traditional bottom-up approaches to infrastructure design focused on providing powerful infrastructure based on requirements and capabilities defined by IT. This approach was effective at addressing discrete business needs. However, a purely technology focused approach suffers from a lack of integration with the future business direction of the organization. For this reason, businesses are moving to a top-down approach, where executive-level decision makers define strategic infrastructure requirements that support the business as a whole. IT then responds to strategic long-term requirements, rather than focusing on near-term issues.

This shift demands systems that can respond to multiple computing needs—a requirement beyond the capabilities of traditional dedicated infrastructure. Ideally, we should be able to repurpose resources in real time to respond to changing business needs.

Virtualization is key for enabling this change, along with other mainstream technologies such as grid computing and grid-enabled enterprise applications. These technologies promise to underpin innovative new architectures that allow us to deploy systems rapidly and control operating costs.

The Promise of Service Orientation

To help businesses respond more quickly and cost-effectively to changing business needs, enterprise architects have advanced service-oriented concepts as a key design modality. This model builds on several generations of industry best practices to provide a top-down design methodology, where business models drive enterprise architecture, rather than the other way around.

Service orientation enables the reuse and interoperability of enterprise resources such as business applications, information, and data. Within this approach, a virtual infrastructure

(VI) provides compute, networking, and storage resources, as shown in Figure 1, that we can reallocate on demand.

This approach facilitates the goal of providing IT as a service, delivering enterprise resources such as applications, processing power, storage, and network connectivity as loosely coupled services that can be freely interconnected to create new pools of resources. If we explicitly design resources that can be connected in any combination, linking them becomes trivial and building new systems becomes extremely cost-effective.

This allows rapid innovation to support changing business needs. For example, suppose a new initiative requires processor-intensive data mining by a line-of-business application using disparate data sources. This requirement could be daunting under traditional topologies. Under an IT service model, however, we could allocate computing resources as needed and connect software components using well-defined, standards-based interfaces.

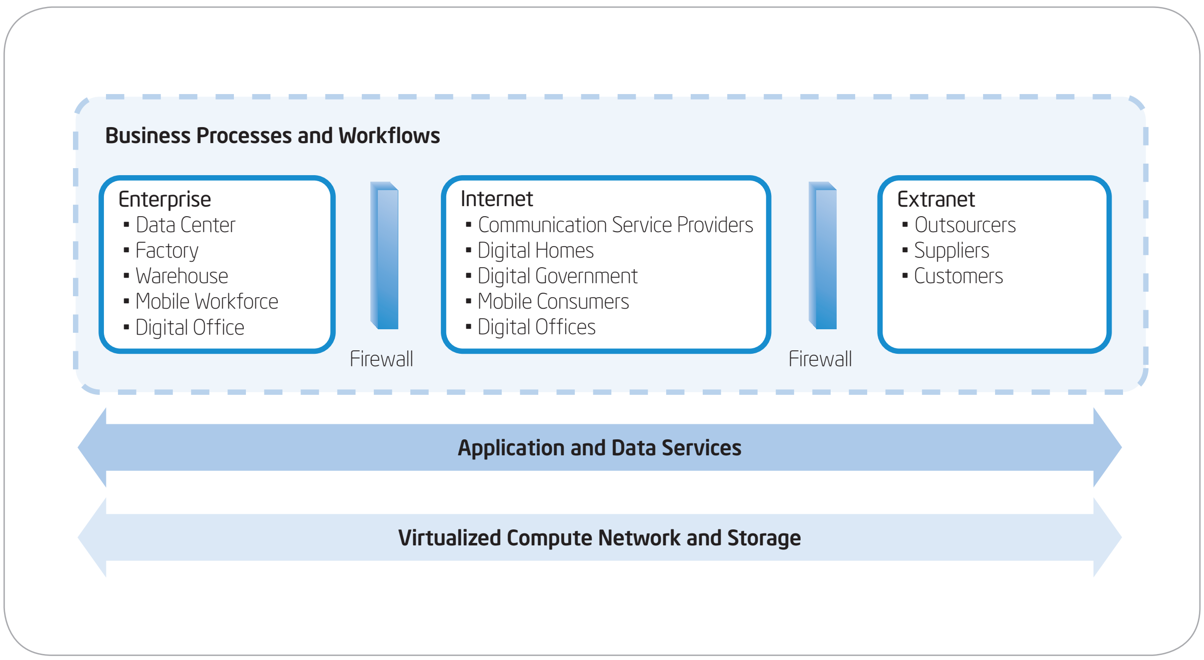

Figure 1. Support for business processes and workflows in a service-oriented environment.

Next-Generation Execution Resources: Virtualized Compute Infrastructure

Intel's vision for the next-generation enterprise includes an unprecedented level of computing resources abstraction. Our service-oriented architecture goal includes a virtualized compute infrastructure (VCI) that will enable us to meter and redirect processing resources according to changing business needs, as governed by IT policies designed for that purpose. This business model also allows us to track and bill resource use to individual users or groups.

By using this VCI, we are seeking to transform dedicated resources into virtualized and pooled ones. Traditional, non-virtualized systems require static, manual allocation and provisioning; VCI makes allocation and provisioning dynamic.

Proof of Concept

Intel created a proof of concept (PoC) VCI implementation in the lab, based on an event-driven architecture that responds to workload demands according to pre-set policies. The PoC consisted of a data-mining system designed to support business intelligence (BI) operations.

Baseline: Non-virtualized Environment

As a baseline for comparison, we created a non-virtualized enterprise computing environment, as shown in Figure 2 on the next page. We based our hypothetical scenario on a user who works for a financial services firm that offers home equity loans to customers nationwide. She needs to analyze large volumes of loan applications created daily at banking branches across the country. The system uses advanced analytics to predict the risk associated with loan applications based on previous loan history, creating scoring models that it uses to rate incoming applications.

To prepare these models, the system pulls loan-application data from a large number of external sources and then transforms, analyzes, and processes the data according to the appropriate business logic. The system then uses reporting tools to present the results to the user.

The main steps, as diagrammed in Figure 2, are:

1. The business user in charge of running the analysis launches a series of jobs using the Web-based client, which sends the appropriate database queries through the portal server.
2. The portal server connects to the grid-based array of BI servers that assemble the data from multiple sources.
3. The BI servers request data from back-end data sources using extraction, transformation, and loading (ETL) servers.

4. Dedicated ETL servers, managed by the platform job scheduler, execute the job.
5. Databases store the transformed data.
6. The system sends transformed data and metadata to various reporting clients through the BI servers and a Web portal.

Because loan applications from each banking branch can be processed in parallel, this problem is well suited to grid computing. The grid topology also enables excellent scalability, because the number of compute nodes can be increased as the number of branches and loan applications grows.

However, a disadvantage of the non-virtualized approach is that adding more compute nodes means purchasing and maintaining a large amount of hardware that is used only during a relatively small percentage of the work day.

Virtualized Compute Infrastructure

Once we had the baseline system up and running, we were ready to ascertain whether we could incorporate a server pool of virtual machines (VMs) into the architecture. This server pool would function like a utility of metered virtualized compute resources that we could apply anywhere in the enterprise.

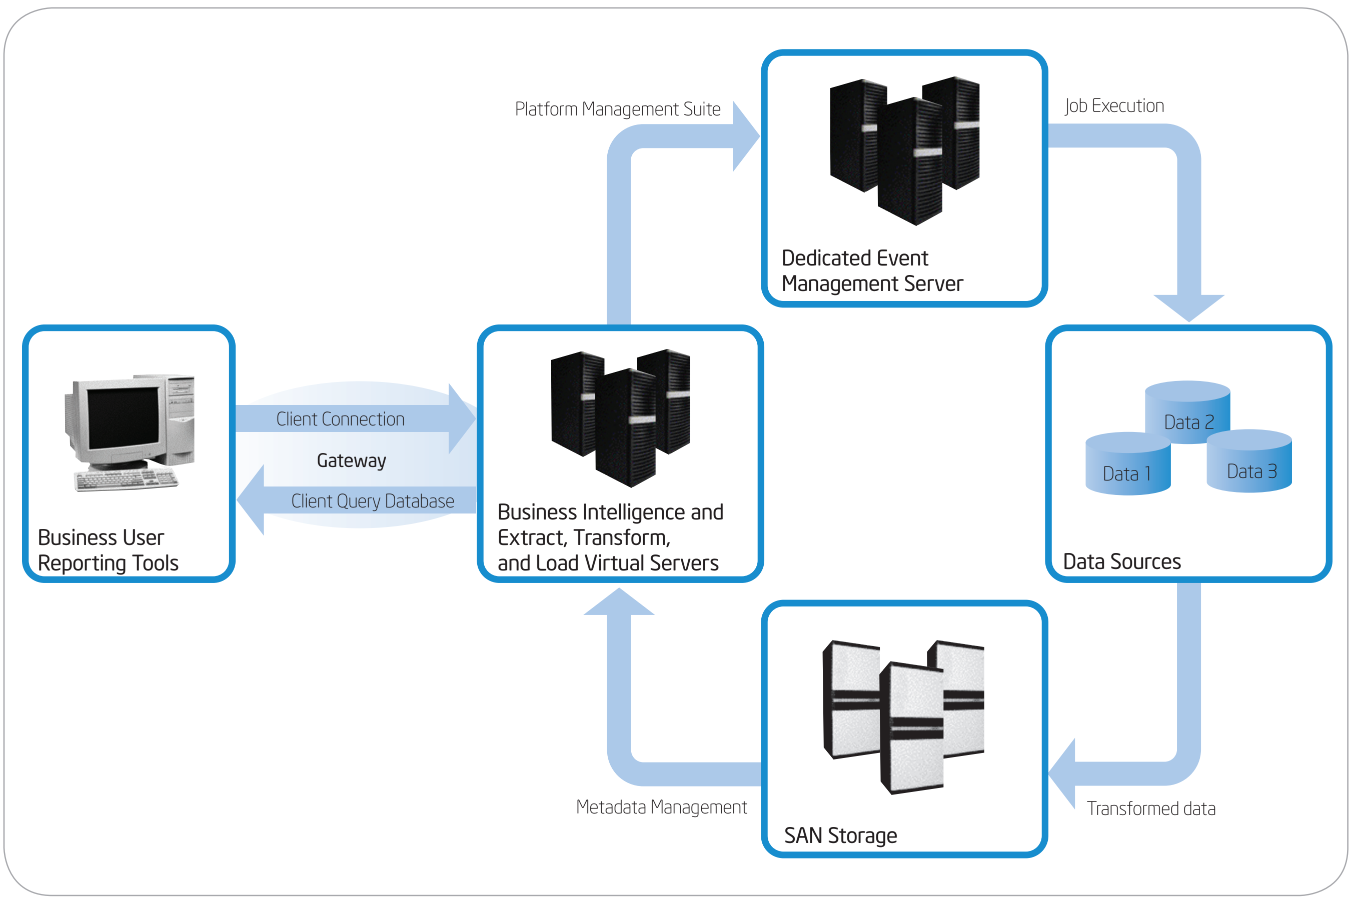

Figure 2. Traditional non-virtualized enterprise computing environment.

In this VCI, a monitoring agent identifies when the four physical nodes are operating at capacity. Once they reach a given threshold, an application that manages VMs activates additional compute resources, adding to the cluster six VMs resident on separate hardware.

The PoC infrastructure consisted of 12 servers, including administrative nodes. In a real-world implementation, any number of systems throughout the enterprise could share these virtualized compute resources.

Results

This PoC primarily aimed to show that we could apply virtualized compute resources to workloads on the basis of pre-defined policies, while preserving the integrity of analysis operations. As demonstrated by the charts in Figure 3, the data mining results produced by the non-virtualized and virtualized environments were identical, confirming that we could preserve the integrity of the analysis.

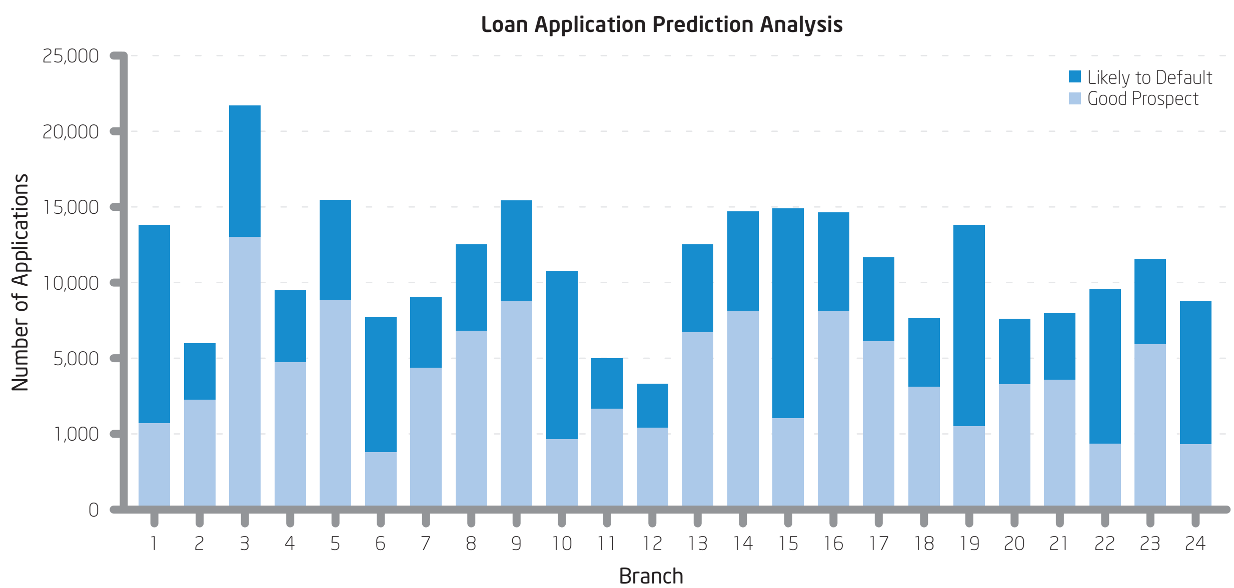

Note: Virtualized and Non-Virtualized Results are Identical

Compute Infrastructure	Nodes Available	Analysis Time
Traditional	4	11.88 minutes
Virtualized	10	7.50 minutes

Figure 3. Loan prediction analysis software ran more quickly in our virtualized environment. We ran test data mining software that analyzed thousands of applications at a fictional financial institution with many branch offices, determining whether each application represented a good prospect or one likely to default. The test completed much more quickly in the virtualized environment than in the non-virtualized environment. The results of the analysis were identical in both environments, confirming that the integrity of the analysis was preserved in our virtualized environment.

However, our virtual infrastructure also produced a dramatic increase in performance. Response time shrank from 11 minutes 53 seconds in the non-virtualized environment to 7 minutes 30 seconds with the VCI, a decrease of approximately 37 percent.

As a result of this testing, we expect that automated, policy-driven management of virtualized compute resources could result in:

- More manageable and efficient resource allocation
- Improved utilization, with a smaller number of servers needed for a given workload
- Decreased risk, with less dependence on a single point of failure
- Excellent investment protection

Expanded Benefits: Orchestrating Multiple Applications

To capture the full benefit of VCI, we have to make the virtualized processing resources available throughout the organization, rather than statically allocated to specific user groups. For instance, in our loan application example, we might add a set of data integration users to the existing BI users. To obtain higher resource utilization, we would share the virtualized compute resources between the two user groups.

As in our PoC, the system would monitor resource utilization on local servers according to preset policies and allocate additional resources from the virtualized compute pool when it detects bottlenecks.

We're also considering more sophisticated models of resource allocation, such as time-based allocation. In the simplest case, we could set policies that allocate compute resources to one group during the day and to a different group at night. This model might be appropriate for batch jobs that are performed nightly or monthly.

We could also give higher priority to mission-critical operations. By combining prioritization with other policy considerations and schedules that vary according to the time of day or day of the week, we could fine-tune the environment to direct virtualized resources where they are most valuable.

We expect to gain significant benefits by deploying virtualized resources throughout the enterprise using policies that respond to events, workflows, and changing business needs. As we fully implement this model, we will monitor performance and identify other business solutions that we can incorporate into the virtualized infrastructure. We will also look for opportunities to take advantage of future developments in hardware platforms, virtualization software, and other components.

Intel is working with multiple vendors to enable further performance and flexibility in the next generation of grid-enabled solutions. Many grid computing solution vendors are working to optimize their product offerings to take advantage of Intel® Virtualization Technology (Intel® VT). Together with current work to improve the robustness of virtualization software, we expect that these advances will help provide additional capacity as we continue to develop our VCI.

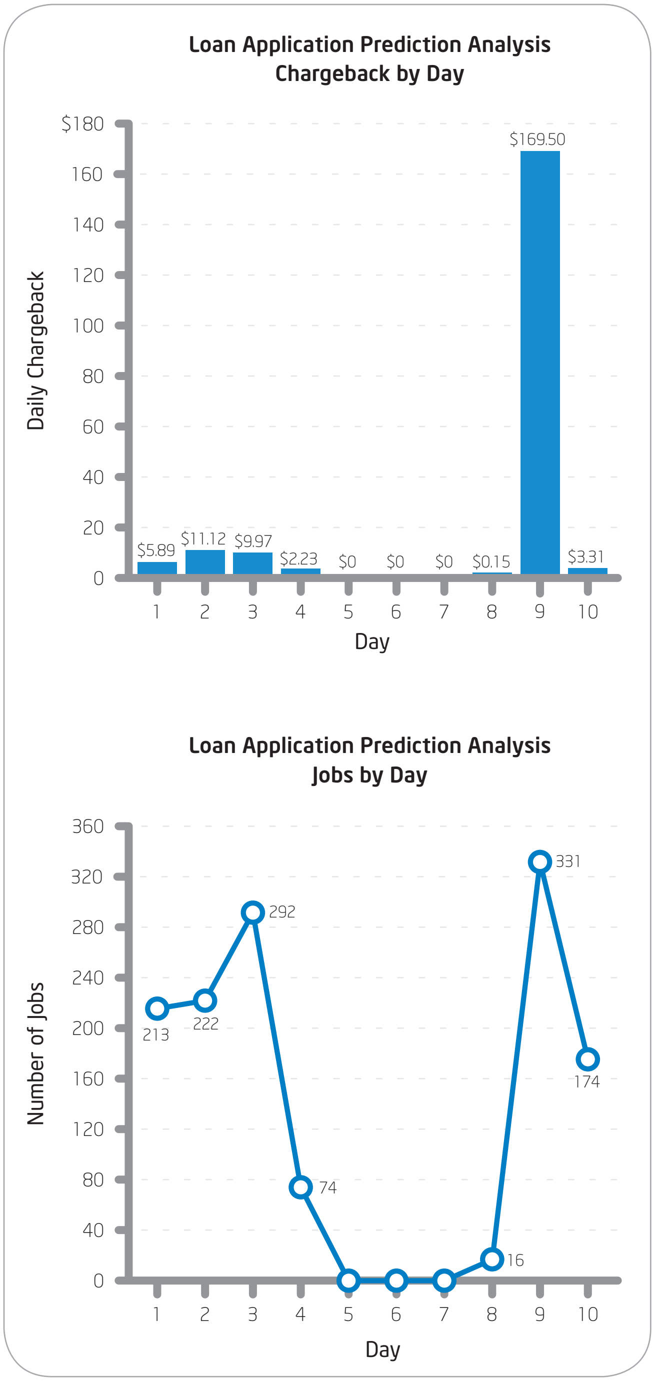

Figure 4. Chargeback reporting in a virtualized environment. We used reporting tools to gather data about how resources were used in our proof of concept VCI. Using this data, we could create utility-like chargeback billing based on predetermined rates per CPU hour, as shown.

Utility Chargeback Capability

Intel IT is also examining the potential to collect and report usage data associated with each business unit. We may eventually use this functionality to support a utility-like billing scheme, with chargeback billing to groups that use virtualized computing resources. This would further Intel's ability to run IT like a business.

Our PoC system supports this mechanism well. We used reporting tools to collect actual data from the running VCI. That data could be used to create chargebacks for individual users or business groups based on predetermined rates per CPU hour, as shown in Figure 4.

We expect that implementing a chargeback scheme will help us measure usage by specific business units and better quantify where IT resources are being used, providing budget tracking to support overall corporate strategy.

Complex Utility Analysis

Using more sophisticated tools, we could move beyond chargeback into more complex modeling and forecasting. We envisage executive portals and dashboards that could deliver visually rich and compelling reports in real time, support performance scorecards identifying utilization of the organization's resources, and support IT governance.

These capabilities could help provide financial transparency for the IT organization, including hypothetical cost modeling and forecasting. They could also provide transparency into the services themselves, helping us streamline processes and analyze latent or chronic performance and stability issues. The capacity forecasting, threshold and peak-period analysis, and workload characterization and profiling capabilities could help Intel IT provide strategic input and reporting to senior executive management

Conclusion

Our successful PoC indicates that a VCI has great potential for improving agility within the IT organization. As a result, Intel IT has begun building a VCI similar to the PoC environment and the outcome of that work so far is very encouraging. This infrastructure will virtualize compute power and distribute it across the production environment to business groups according to their shifting needs.

The hardware foundation for our virtualized architecture is the Intel® architecture-based service-oriented infrastructure (SOI), including platforms and technologies such as Intel VT, and Intel® Input/Output Acceleration Technology (Intel® I/OAT).

Based on preliminary analysis, we expect the following business advantages:

- **Lower total cost of ownership.** By making better use of compute resources, Intel IT expects to reduce the number of systems required to support equivalent enterprise workloads. Fewer systems should result in lower acquisition, maintenance, and support costs.
- **Improved systems integration among business groups.** By providing enterprise-wide resources that can we can utilize in any solution, Intel IT plans to extend existing initiatives aimed at creating of reusable components in a service-oriented modality.
- **Lower cost of entry for emerging initiatives.** New projects with extensive compute needs will be able to obtain resources with minimal initial investment.

A VCI extends previous agility benefits associated with service-oriented modalities, with the potential to dramatically improve utilization and reduce costs through advanced resource sharing across the business.

VCI is part of a much larger initiative at Intel and within the technology industry to provide IT as a set of services that form a utility-like structure for IT resources. We expect sophisticated chargeback and analysis mechanisms for these resources will increase operational and financial transparency of the system, allowing IT to be more responsive to the strategic needs of the company as a whole.

Authors

Parviz Peiravi is a senior solutions architect with Intel's Customer Solutions Group.

Acronyms

BI	business intelligence
ETL	extraction, transformation, and loading
Intel I/OAT	Intel Input/Output Technology
Intel VT	Intel Virtualization Technology
PoC	proof of concept
SOI	service-oriented infrastructure
VCI	virtualized compute infrastructure
VI	virtual infrastructure
VM	virtual machine

www.intel.com/IT

Printed in USA ♻ Please Recycle
1006/ARM/RDA/PDF Order Number: 315021-001US