

White Paper
Intel Information Technology
Computer Manufacturing
Enterprise Resource Planning

A Quantitative Approach to ERP Server Platform Sizing

Under-sizing server platforms for an environment may lead to mid-life refresh, resulting in significant disruption to business operations, added logistical complexity, and an increase in total cost of ownership. Intel IT developed a quantitative model to assist with server sizing by analyzing the effects of many different factors on enterprise resource planning (ERP) server utilization. Our analysis shows that the headroom advantage offered by four-socket servers can increase flexibility and help avoid mid-life refresh in situations where there is greater uncertainty in workload growth forecasts and larger demand peaks. This headroom is also important when targeting longer server refresh cycles of four years or more and when using active-active cluster designs to provide high availability.

Sudip Chahal and Karl Mailman, Intel Corporation

July 2009

IT@Intel

Executive Summary

Server sizing and selection is a critical element of Intel IT's enterprise resource planning (ERP) infrastructure strategy, which makes use of both two-socket and four-socket servers. Under-sizing servers may result in the need for mid-life server refresh, which can result in significant disruption to business operations, added logistical complexity, and increased total cost of ownership.

We developed a quantitative model to assist with server sizing by analyzing the effects of various factors on the utilization of the primary server resources: compute, memory, and I/O.

We developed a quantitative model to assist with server sizing by analyzing the effects of various factors on the utilization of the primary server resources: compute, memory, and I/O. The model is based on the Intel IT approach to sizing servers for our ERP environment.

Our analysis shows that four-socket servers, which provide greater headroom than two-socket servers, may be preferable in a number of situations:

- When there is greater uncertainty in workload growth forecasts, including larger workload peaks.
- With longer refresh cycles: Four-socket servers can accommodate greater workload growth over a longer period of time. The high logistical complexity of an upgrade makes it very desirable to plan for a refresh cycle of four or more years and to avoid mid-life refresh.
- With active-active cluster designs, which necessitate planning for greater headroom.

We need to consider the combined effect of multiple factors when determining whether a server has enough headroom to support mission-critical ERP applications without requiring a mid-life refresh. Based on our evaluation of these factors, Intel IT uses four-socket servers for our most demanding ERP instances.

Contents

Executive Summary 2

Background 3

Server Sizing Approach 4

 Factors Impacting Server Platform Selection and Sizing 4

 Modeling ERP Scenarios 6

 Risks of Under-Sizing 8

 Intel IT ERP Server Strategy 11

Conclusion 11

Authors 12

Acronyms 12

Background

Intel IT operates a very large decentralized enterprise resource planning (ERP) environment based on industry-standard two-socket and four-socket servers. We have about 10,000 active users, and our environment includes many ERP instances running on approximately 250 servers. We have found that a decentralized ERP approach offers several advantages, including lower server acquisition costs and increased flexibility and agility.

When we select a server platform to run a specific ERP instance, our goal is to size the server so that we provide predictable high performance for this critical environment and minimize disruptions to the business due to changes to the ERP environment during the typical Intel IT four-year refresh cycle. At the same time, we seek to minimize server total cost of ownership (TCO).

Various factors determine our ERP platform selection and sizing decisions, including workload growth projections, maximum utilization target, the size of workload spikes, and the capacity of two-socket and four-socket servers.

We have developed a quantitative model that assists with server platform selection and sizing by enabling us to analyze the effects of these factors on platform utilization.

Server Sizing Approach

We consider a range of factors that significantly affect utilization of the key server resource categories: compute, memory, and I/O. We model the combined effect of these factors to understand the overall implications for server sizing and selection. We use this approach to make sure that we select servers with the right amount of headroom to support the expected demands of the workload over the server's planned four-year life.

Factors Impacting Server Platform Selection and Sizing

A variety of factors impact server platform selection and sizing:

- Initial measured average utilization
- Maximum utilization target
- Peaks in demand
- Workload growth projections
- Relative capacities of two-socket and four-socket servers
- Server refresh cycle
- Advanced considerations such as clustering and failover

Any of these may significantly affect platform utilization and selection. To illustrate the concepts, we will examine many of these factors individually. However, when sizing and selecting servers, it is essential to consider the combined effect of all the factors.

Initial Measured Average Utilization

Initial measured average utilization provides the starting point for assessing the required server capacity. We assess the effects of all other factors—such as workload growth rate or the size of demand peaks—relative to this initial, baseline utilization level.

Maximum Utilization Target

For a business-critical ERP environment, it is essential to maintain good responsiveness under all conditions.

To achieve this, we need to establish a maximum utilization target for key server resources. This target depends on factors such as the specific business requirements and the resource type. If this utilization target is exceeded, response times may increase dramatically, endangering the ability to meet service-level agreements (SLAs).

The maximum utilization target for the key server resources—processor and memory—typically ranges from about 65 percent to 80 percent. For illustration purposes, in this paper we assume a maximum utilization target of 75 percent.

Examples of Workload Peaks

Short-term variation in demand can drive significant swings in resource utilization that last from a few minutes to more than 24 hours. Examples we have experienced include:

- **Requests to employee intranet portal.** In response to a message from Intel's CEO about changes to employee stock plans, traffic on the employees' intranet home page surged to a peak of 80,000 requests per hour—8x the average.
- **User-initiated execution of business processes.** When users initiate business processes, they may greatly increase server compute demand. User-initiated execution of one master data synchronization routine requires processing more than 40,000 records—an 8x increase over the standard batch process.
- **Reconciling inventory.** Each week, the task of reconciling and valuing the inventory at distribution centers worldwide creates a six-hour surge in processor utilization to about 70 percent—roughly 3x the average of 20 to 25 percent utilization.
- **Updating test systems.** As a production system database grows, the task of copying code and data to test systems becomes increasingly resource-intensive. One quarterly process involves updating each row of a table with many millions of entries; this process alone takes more than 24 hours and consumes 20 to 25 percent of the server's processing capacity.

Peaks in Demand

Short-term bursts in demand can drive significant swings in resource utilization (see sidebar). These bursty workloads can result in utilization peaks that are much higher than the long-term average. In general, we have noted that for many services the peak workload exceeds the average by 2x to 10x.

We base maximum utilization targets on peaks that are sustained over a few minutes, rather than on instantaneous peaks. We have found that peaks are typically more pronounced for processor and I/O utilization than for memory utilization. Our experience is that processor utilization can surge up to 10x or higher during a short period, but peaks in memory utilization are more likely to be in the range of 1.3x to 1.5x. Relative to long-term averages, the peak I/O utilization can vary to an even greater extent.

When planning server capacity, we account for peak utilization across each key resource to make sure it does not exceed our maximum utilization target during the life of the server, otherwise we risk reduced service levels to our IT customers.

Workload Growth Projections

The average utilization generated by a workload typically grows over the life of a server. Factors that affect this growth include:

- Anticipated growth in the number of users
- The expected addition of new applications on the server

- OS and application refreshes; for example, migration to a unicode database can result in significant increases in processor utilization and database size
- Changes in average transaction complexity and transaction resource requirements

Although the resulting aggregate workload growth may vary from year to year, it is useful to model growth in average utilization as compound annual growth rate (CAGR). We have found that CAGR for ERP workloads typically ranges from 5 to 30 percent. For illustration purposes, in this analysis we will model both 10 percent and 20 percent CAGRs.

Relative Capacities of Two-socket and Four-socket Servers

The decision whether to select a two-socket or four-socket server depends on the relative capacities of the servers. The relative capacity varies over time as new server platforms are introduced, and depends on the resource—compute, memory, or I/O—being compared. General capacities are shown in Table 1.

Server Refresh Cycle

Server sizing is significantly affected by an organization's server refresh cycle. Servers must continue to offer good responsiveness throughout their life, so they must be sized to accommodate the expected workload growth over the entire refresh cycle. The longer the planned refresh cycle, the more a workload is likely to grow over the life of the server, resulting in a need to plan for additional headroom when initially sizing the server.

Table 1. Relative Capacities of Two-socket and Four-socket Enterprise Resource Planning Server Platforms

	Two-socket Server	Four-socket Server	Benefits of Four-socket Servers over Two-socket Servers
Number of cores	4 to 8	16 to 24	<ul style="list-style-type: none"> ▪ Performance scaling: typically 1.5x to 2x ▪ Range of 1.1x to 2.2x observed
I/O slots	Typically 2 to 5	Typically 6 to 8	<ul style="list-style-type: none"> ▪ Significantly greater I/O expandability: typically in the 2x range
Memory slots	18	32	<ul style="list-style-type: none"> ▪ Greater memory capacity: approximately 2x to 4x
Maximum memory capacity	72 GB with 4-GB DIMMs	128 GB with 4-GB DIMMs	

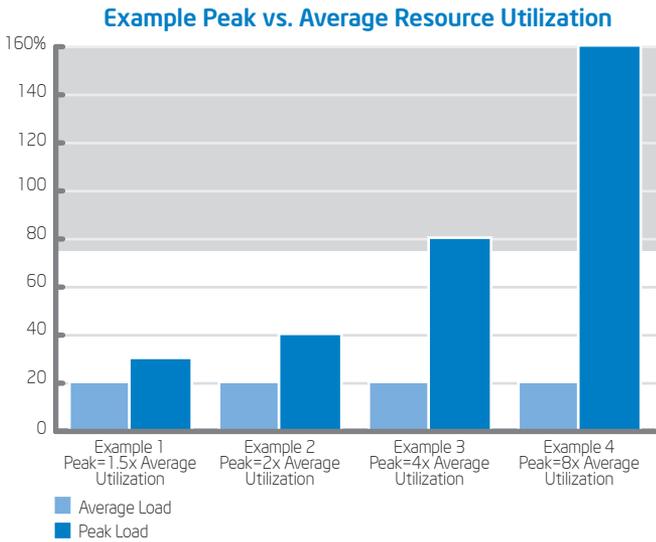


Figure 1. Modeling the impact of workload peaks on server utilization.

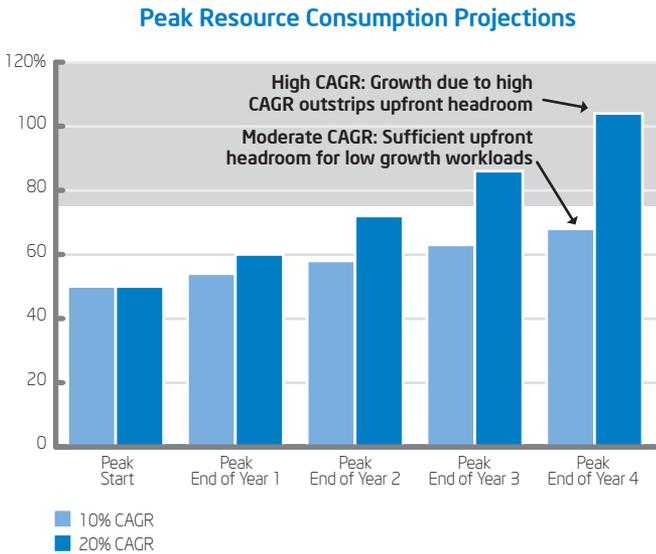


Figure 2. Modeling the effect of different workload growth rates on server utilization. We assume a four-year hardware refresh cycle and illustrate both 10 and 20 percent compound annual growth rates (CAGRs).

Advanced Considerations: Clustering and Failover

We use clustering to provide high availability for critical ERP instances. One of these clustered instances, the production instance, runs our production ERP transactions. Similar clusters are used for benchmark testing, using two-node active-active clusters, which provide maximum performance and headroom during normal operations.

When failover occurs, the surviving server must run the combined workloads of both original servers, and continue to deliver good performance to meet our SLAs. This means that we need to size each server so that it can support the combined workloads of both servers in the cluster, without exceeding its maximum utilization target under anticipated peak loads.

Modeling ERP Scenarios

We have developed a model that shows how variation in these individual factors affects utilization of key resources over a server’s refresh cycle. The resulting predicted utilization levels determine whether we should select a two-socket or a four-socket server.

All assumptions used in the following modeled scenarios are for illustration purposes only. We assume a maximum utilization target of 75 percent. Actual workload growth projections, peak and average utilization ratios, utilization levels, and relative server capacities will depend on the server resource (processor, memory, or I/O) and the specific deployment scenario being modeled. Relative server capacities also vary over time as new technologies are introduced in different platforms at different times.

Impact of Resource Utilization Peaks

The size of workload peaks, relative to average utilization, significantly affects server selection. This example is based on a workload that generates 20 percent average utilization, as shown in Figure 1.

The server can accommodate peaks that generate 1.5x average utilization (example 1 in Figure 1). However, a bursty workload that generates peaks 4x the average (example 3) exceeds the maximum utilization target, likely requiring a server with greater capacity; a workload with peaks 8x the average (example 4) would definitely require a higher-capacity server.

Different Workload Growth Rates

This example is based on a workload that initially generates peak utilization of 50 percent. As shown in Figure 2, the server can accommodate an average 10 percent growth rate over the four-year cycle. However, if the workload grows at an average 20 percent a year, it approaches the maximum utilization target in year two, and exceeds the target by the end of year three—well before the planned refresh date.

Server Refresh Cycle

Longer refresh cycles require more headroom to accommodate growth. This example is based on a server running a workload that initially generates peak utilization of 40 percent. With a short three-year refresh cycle, the server can accommodate even a rapid 20 percent growth rate without exceeding its maximum target utilization, as shown in Figure 3.

However, Intel IT generally uses a four-year refresh cycle across our environment, as described earlier. With refresh cycles of four years or more, 20 percent annual growth in this workload would exceed this server’s maximum utilization target; a server with greater capacity would be required.

Clustering and Failover

Clustering is essential to provide high availability for our production instances. This effectively doubles the required capacity of each server, because if one server fails, the surviving server has to run the workloads of both and do so with good performance to meet SLAs. Therefore, in normal operations, peak utilization of the servers in the cluster should be less than half the planned maximum utilization target; otherwise failover can cause utilization on the surviving server to exceed the maximum utilization target, endangering SLAs. If our projections indicate frequent loading above this level, higher-capacity servers are needed to guarantee the required responsiveness.

Sizing a Server Platform for a Large ERP Production Instance

When we size an ERP server platform, we must consider the potential combined effect of all the above factors.

In this example, we examine how these combined factors affect utilization of a key resource—memory—and the resulting impact on selecting a server to run a large ERP production instance.

Our production instances are clustered to provide high availability; each server must be sized so that, in the event of failover, it can run the workload of both servers in the cluster.

This scenario is shown in Figure 4. Our example is based on a workload that initially utilizes approximately 20 percent of the memory capacity of a two-socket server. The same workload initially utilizes approximately 10 percent of the memory capacity of a four-socket server, because the four-socket server has about 2x the memory capacity of a two-socket server.

For an active-active two-node cluster of two-socket servers with an initial cluster utilization of 20 percent, even a moderate 10-percent growth rate combined with peaks 50 percent higher than average can cause peak utilization in failover mode that will exceed 75 percent by the second year. In contrast, the four-socket server-

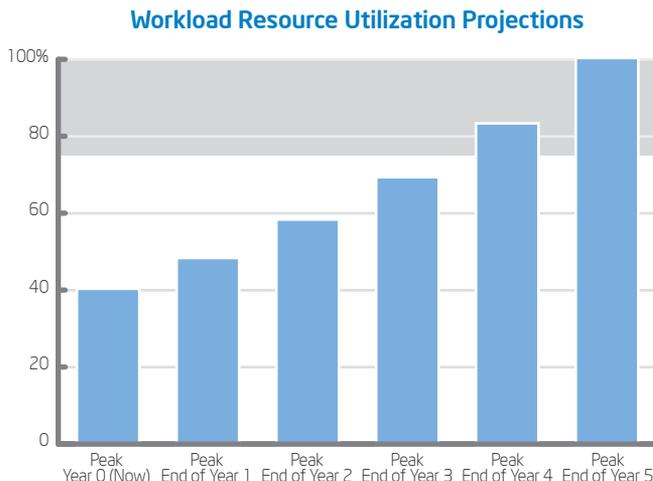


Figure 3. Modeling the effect of different refresh cycles. We assume a compound annual growth rate (CAGR) of 20 percent. With a short three-year refresh cycle, a server platform with modest headroom may suffice. For longer four- to five-year refresh cycles, high-CAGR workloads may necessitate a server with greater capacity.

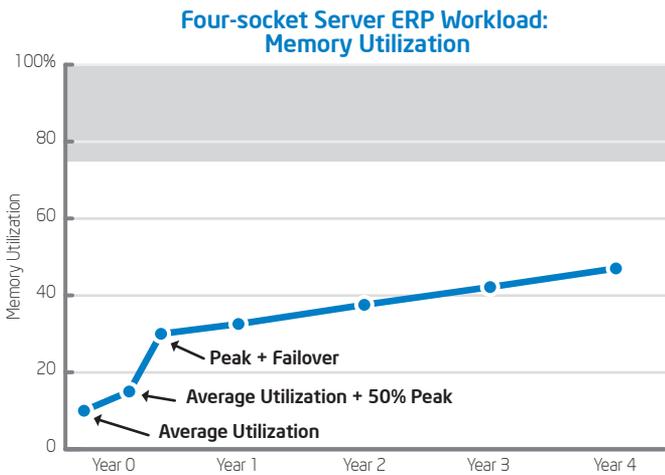
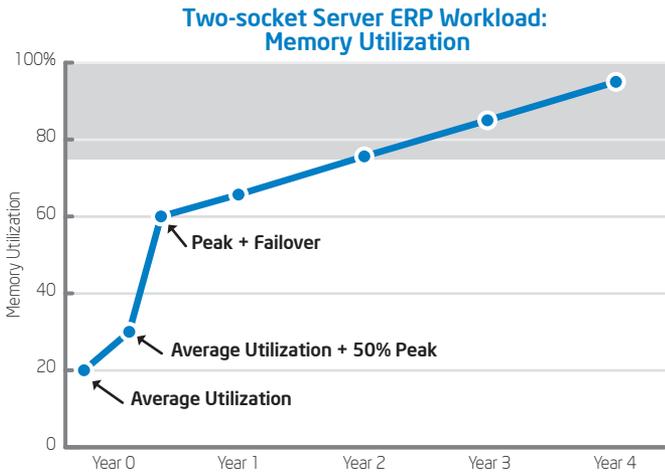


Figure 4. The effect of multiple factors on server sizing for a large enterprise resource planning (ERP) production instance.

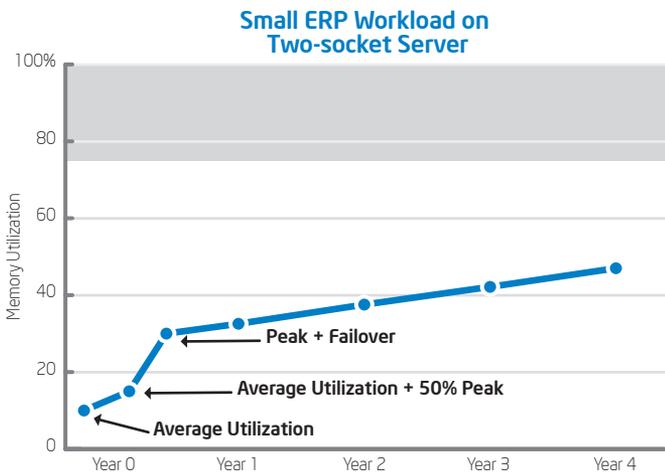


Figure 5. Server sizing for a small enterprise resource planning (ERP) production instance.

based cluster comfortably accommodates growth over our entire four-year refresh cycle.

For smaller workloads, a cluster of two-socket servers may have adequate capacity to accommodate growth over the entire refresh cycle, as shown in Figure 5.

Risks of Under-Sizing

Under-sizing—selecting a server that exceeds its maximum utilization target before the end of its planned lifecycle—has significant consequences.

Each Intel ERP production instance includes a scale-up database and several application components. If the ERP server cannot accommodate growth in the database, the server typically must be refreshed in mid-life in order to help ensure that ERP performance continues to meet SLAs. This mid-life refresh causes significant disruption to business operations, adds logistical complexity and increases TCO.

The following analysis examines the impact of under-sizing, the situation shown in Figure 4, due to selecting a server that is not able to accommodate workload growth over a planned four-year refresh cycle, resulting in the need for mid-life refresh. It compares this approach to selecting an appropriately sized server that has enough headroom for the entire four-year cycle.

TCO Analysis

Mid-life server refresh results in significant hardware and non-hardware costs, due to the need to purchase new servers and to implement and test changes across this business-critical environment.

Our analysis of the effect on TCO is based on published server list prices and estimated industry compensation rates.¹ We included the following factors:

Labor costs

Mid-life refresh is resource-intensive, requiring work by IT and business groups in three main areas. The effort required in each area, and the resulting cost, varies depending on the complexity of the change.

Qualification. We need to qualify each new server platform for use in our ERP environment. Intel IT has a team of three people who work part-time on this, representing a total of about 6 to 12 person-weeks.

¹ www.dice.com

Implementation. A pipeline architect, a database administrator, and an ERP application specialist work together to refresh the servers running each of the ERP instances. This represents a total of about 12 to 36 person-weeks of effort.

Business group regression testing and checkout.

This is the biggest cost, accounting for an estimated 36 to 72 person-weeks of work. Because ERP is so critical to each of Intel's business groups, they tend to be risk-averse and require repeated testing of any changes. They commit resources to regression testing the changes to each instance, as well as doing a complete checkout of the system once we have completed replacement of the servers and software. The effort and cost are multiplied by the fact that each instance is typically shared by many projects; at least one representative from each project is required to perform these tasks across all the business group's ERP instances.

Hardware costs

The hardware costs of a mid-life refresh can vary significantly.

For each Intel business group ERP implementation, there is a pipeline of ERP application instances. Each instance supports a specific function, such as development, benchmark test, production, or production support, and may be implemented using one or more servers.

The hardware costs of a mid-life refresh are determined by the number and type of servers that we need to replace within this pipeline. Several factors can influence this, including:

Differences between platforms. The hardware and software platforms available as replacements may differ considerably from the platforms that were originally deployed, depending on the length of time since the original deployment. If there has been too much change to the hardware, OS, and drivers, it may not be feasible to mix new platforms with the original platforms within a pipeline.

Need for consistency. It is highly desirable, for support reasons, to have as much hardware and software

consistency across an entire pipeline as is economically feasible. This consideration may drive wholesale replacement across a pipeline even if it is conceptually possible to mix new and existing platforms.

Mid-life refresh options

As a result of these factors, there may be a variety of mid-life refresh options to consider. For illustration purposes, using the example of a pipeline originally based on 16 two-socket servers, these options may include:

Replace all 16 servers with new two-socket servers (greatest complexity). This wholesale replacement is the most complex refresh to implement, but may be the best option from a supportability standpoint if there have been substantial changes in the platforms—such as a new OS revision and drivers—so that it becomes impractical to mix new and original platforms within a pipeline.

Replace six servers with new two-socket servers based on a newer processor generation (moderate complexity). In this example, we use new two-socket servers based on a newer generation of processors. We assume that these replacement servers offer enough additional capacity, and that there are minimal changes to the software stack so that it is feasible for the new servers to coexist with the original servers in the same pipeline. In this scenario we limit the scope of replacements to the more demanding instances within the pipeline, which are the production, disaster recovery, and benchmark test instances. We would need to upgrade the latter two instances because they are typically configured to match the production environment.

Replace six servers with new four-socket servers based on the same processor generation (least complexity). In a variation of the preceding example, we use new four-socket servers, assuming that these replacement servers offer enough additional capacity and that there are minimal changes to the software stack. As in the preceding option, the scope of replacements is limited to the more demanding instances: production, disaster recovery, and benchmark test.

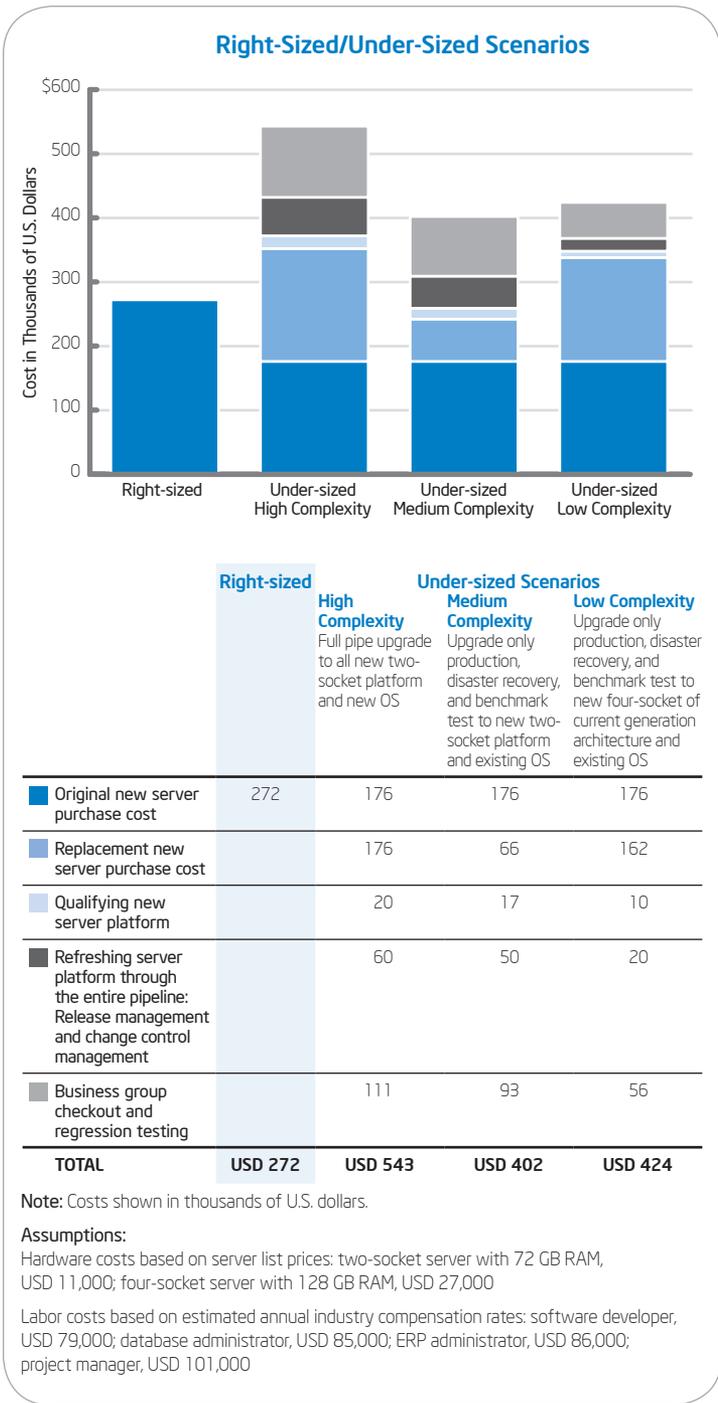


Figure 6. Total cost of ownership analysis of mid-life server refresh options.

Total cost

Figure 6 illustrates estimated total cost for the examples above. As shown, all the mid-life upgrade options result in substantially increased TCO. Depending on the new hardware required, TCO with a mid-life upgrade ranges from approximately 1.5x to 2x the cost of a “right-sized” approach based on originally selecting four-socket servers with enough headroom for the entire four-year refresh cycle.

Business Disruption

The greatest impact of mid-life refresh—even more significant than the cost of acquiring new servers—is the business disruption as new instances are tested and deployed on the new servers.

There are many impacts. These include the need to allocate testing and validation resources, as described above, that could otherwise have been applied to other projects. Other impacts include the downtime required on each system during the implementation of each change.

In addition, production systems are unavailable for use during the cutover to the new servers. This results in idle resources and potentially the need for extra work after the cutover to process anything that was missed.

The high logistical complexity of an upgrade makes it very desirable to plan for a refresh cycle of four years or more and avoid mid-life upgrades wherever possible. Business groups are extremely concerned about changes to the critical ERP environment—much more so than for other environments that have less stringent requirements and correspondingly lower barriers to change. In the ERP environment, each additional IT request for resources and downtime may face increasing scrutiny and resistance.

Intel IT ERP Server Strategy

Our analysis has resulted in the current Intel IT ERP server positioning framework, shown in Figure 7.

If peak requirements of an ERP database workload are expected to frequently exceed 40 percent of the processor or memory resources of a two-socket server cluster within the four-year refresh cycle, we standardize on a four-socket server cluster. This is because in a failover situation, the surviving server must be able to run the entire workload that was previously supported by both servers and continue to deliver good performance while doing so.

As a result, we use four-socket servers for our most demanding production, benchmark, and disaster recovery instances. This approach enables us to avoid the cost impact and disruption caused by mid-life server refresh.

We use two-socket servers for many non-production roles as well as smaller production instances.

Over time, we expect that four-socket servers may support additional roles as virtualization hosts. We anticipate steady adoption of virtualization in the ERP environment, starting with the lowest-risk instances and progressing to more critical server roles as virtualization delivers proven stability, reliability, and performance. Four-socket servers offer advantages as virtualization hosts because of their greater memory headroom and larger number of processor cores.

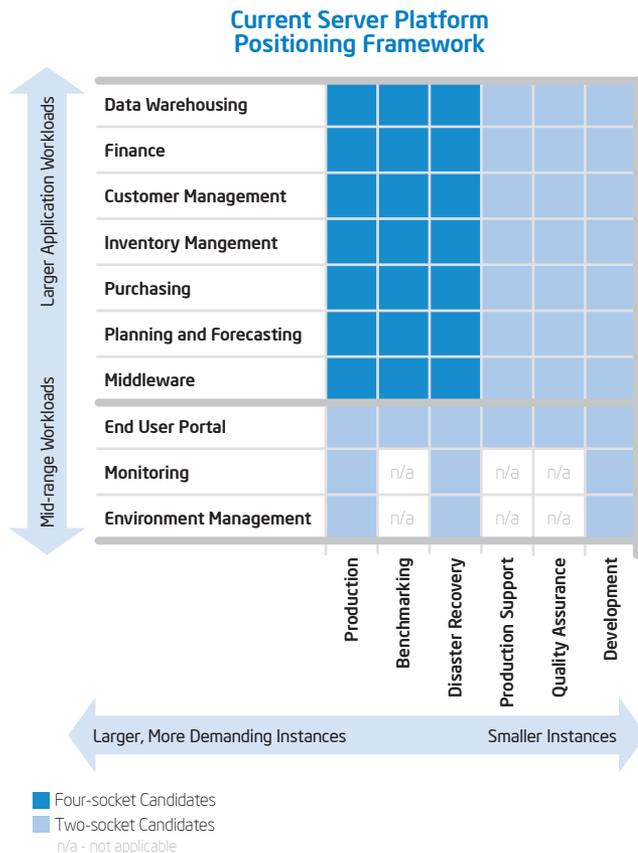


Figure 7. Current Intel IT server platform positioning framework.

Conclusion

Intel’s ERP strategy has enabled Intel IT to effectively support more than 10,000 active users with industry-standard servers.

Server sizing and selection is an essential part of this strategy. Under-sizing servers leads to mid-life refresh, which causes significant business disruption and substantially increases TCO.

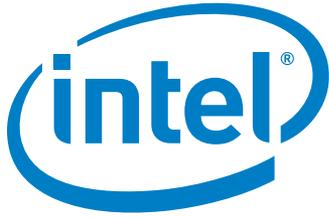
Our quantitative analysis shows that four-socket servers, which provide greater headroom than two-socket servers, may be preferable in the following situations:

- When there is greater uncertainty in the workload growth forecasts, including larger workload peaks.
- With longer refresh cycles: Four-socket servers can accommodate greater workload growth over a longer

period. The high logistical complexity of an upgrade makes it very desirable to plan for a refresh strategy of four or more years and avoid mid-life upgrades.

- With active-active cluster designs, which necessitate planning for greater headroom.

We need to consider the combined effect of these factors when determining which server has enough headroom to support the critical ERP environment without requiring a mid-life refresh. Based on our evaluation of these factors, we use four-socket servers for our demanding most ERP instances.



Authors

Sudip Chahal is a principal engineer with Intel IT.

Karl Mailman is an ERP infrastructure architect with Intel IT.

Acronyms

CAGR compound annual growth rate

ERP enterprise resource planning

SLA service-level agreement

TCO total cost of ownership

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 and the Intel logo 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 © 2009 Intel Corporation. All rights reserved.

Printed in USA
0709/KAR/KC/PDF

 Please Recycle
321374-001US