WHITE PAPER
Intel® Server Platforms
Micro Server Technology

Flexible Scale-Out with Lightweight Micro Servers

Large numbers of small compute nodes based on Intel® architecture can address emerging workload needs and achieve new levels of efficiency and density.

The micro server represents a new server architecture characterized by many lightweight server nodes bundled together in a shared chassis infrastructure. This topology is designed specifically for density, lower power per node, reduced costs, and increased operational efficiency. By sharing common fans, switching, power supplies, and the metal chassis, micro servers streamline rack infrastructure and improve density over standard rack options.

Micro servers have emerged as an option for specific classes of workloads that tend to scale well with large numbers of relatively lightweight nodes. Intel supports this topology with a broad range of processor offerings engineered to give flexibility of choice, improve node density, and reduce node power consumption. This paper introduces the design concepts behind micro servers and gives decision makers the background to make informed decisions about whether micro servers can help them address their server infrastructure needs.

The paper explains the context in which micro servers have evolved and the characteristics of the new design modality. It offers guidelines for identifying whether workloads are well suited to micro server deployments. The paper also provides an overview of Intel® platforms associated with the micro server option and explains how that hardware fits within the framework of industry standards and product offerings.

Overview

As the computing industry innovates, new applications and software models are emerging rapidly—particularly those in the Internet data center. For example, the web tier has grown to become a large portion of the server infrastructure in many Internet data centers, handling content for millions of users. Highly parallelized workloads such as analytics and non-relational databases are becoming more widespread. And a host of new content delivery services are driving different ways to store and distribute digital content.

While the majority of emerging Internet workloads requires high performance systems to handle dynamic, complex tasks, a subset places lower demands on server infrastructure. These lightweight applications don't take full advantage of mainstream server platforms that are richly configured with memory, disk, and networking options typically found in enterprise data centers today.

Micro servers are an emerging category of system designs that have been created to address lightweight applications and simultaneously improve data center efficiency. Micro servers are characterized by large numbers of server nodes configured to share infrastructure such as power and cooling fans in a common chassis.

The trend toward density isn't new. Innovations in server architecture over the past 15 years have helped drive improved efficiency, density, and manageability. Blade servers, for instance, have been widely adopted in the enterprise, and more recent system designs such as half-width rack-mount servers have primarily targeted Internet data center implementations. As shown in Figure 1, micro servers are now taking their place among other server form factors as a new density-optimized option.

April 2011

Expanding Compute Flexibility: Gauging the Micro Server Option

A key ongoing challenge for computing professionals is to identify where their requirements fit along a spectrum of small-scale to large-scale systems. With the micro server category, a new set of options exists for architects as they consider whether a large number of relatively modest nodes is preferable to a smaller number of more powerful ones for a specific implementation.

Distributed, highly parallel workloads with relatively low compute requirements per node can be scaled out effectively in some cases using micro servers. While this class of system is not suited to every workload, it holds particular promise for some market segments. For example, a hosting environment often must support very large numbers of computationally light tasks, such as requests for login authentication or serving small, static HTML pages.

The large number of small nodes in micro server topologies can be ideal for such jobs, as the software is highly parallel and transactions are easily distributed into smaller computational tasks. However, the more computationally intense tasks that make up the majority of jobs in the enterprise are likely to overwhelm small-scale micro server nodes' resources, such as compute power, memory footprint, networking, and storage. Traditional server infrastructures are more effective in those cases.

The diagram in Figure 2 provides a means to more closely consider the suitability of micro servers for various server workloads. In this figure, each of three important server attributes is mapped to a specific axis to generate the visualization, and server workloads can be plotted according to each axis to place them within the model. The three axes are as follows:

Innovation in Server Form Factors

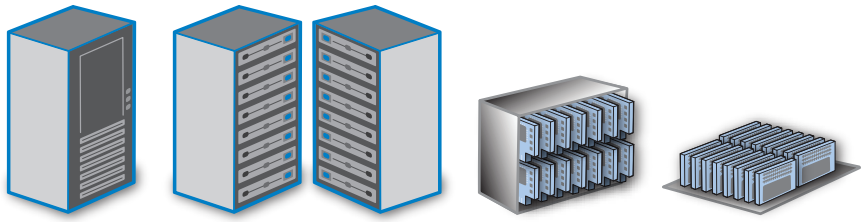

Figure 1. Micro servers have emerged as a new density-optimized category that will co-exist with tower, rack, and blade form factors.

- **Scales with core count**. Higher values on this axis (as they progress to the left) correspond to highly parallel workloads that increasingly benefit from more cores being added to an individual server, so that each unit of work shares common I/O, system memory, storage, and so on.
- **Scales with physical nodes**. Higher values on this axis (as they progress to the right) also correspond to highly parallel workloads, but they tend to scale better by adding more physical servers because of a tendency toward lower memory, I/O, and storage requirements per server.
- **Scales with brawny cores**. Higher values on this axis (as they progress upward) scale best with robust cores, in areas such as frequency, cache, and memory bandwidth. These workloads tend to be single-threaded or have limited parallelization.

The area in Figure 2 identified as "Workloads Suited for Micro Servers" represents workloads that scale best by adding large numbers of lightweight servers. Examples of workloads in this area are ultra-low-end web serving, low-demand (static) content delivery, and basic dedicated hosting.

In practice, only a limited subset of server workloads is well suited to scaling through massive parallelization among many lightweight nodes. The micro server approach is not suited to workloads in many segments, such as high performance computing, financial services, virtualized infrastructure, mission-critical computing, and databases. Two-way or four-way servers based on the Intel® Xeon® processor continue to be the optimal choice for efficient performance on these mainstream workloads.

Matching Capability to the Task: A Range of Intel® Platform Options

Even within the micro server category, there is no one-size-fits-all answer to system design or processor choice. Some micro servers may have impressive brawny single-socket processors with robust memory and storage, while others may have a far higher number of miniature dense configurations with lower power and relatively lower compute capacity per node. To meet the full breadth of these requirements, Intel provides a range of processor options from the smaller to the larger end of the spectrum.

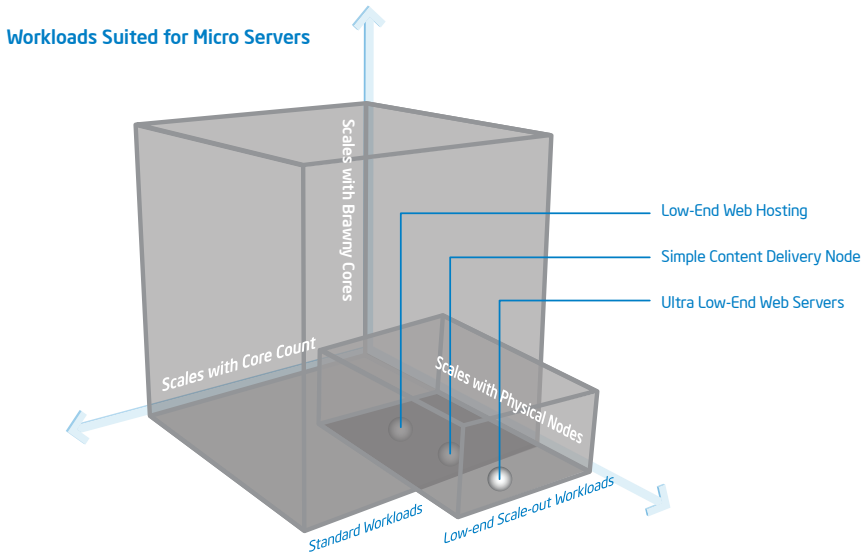

Figure 2. This 3D diagram shows that ideal workloads for micro servers are those that tend to have relatively lower processing requirements and scale best by adding more interconnected physical nodes.

Intel Xeon processors are well suited to applications in many segments, such as high performance computing, financial services, cloud computing, and databases. For the vast majority of micro server workloads, the Intel® Xeon® processor E3 family will be the "sweet spot" for performance, performance per watt, and flexibility. For a limited set of highly parallel workloads where system density is more important than node performance, customers may want to explore the use of the low-voltage Intel Xeon processor E3 family SKUs.

Because all these platforms share a common instruction set architecture, application code can run across the full spectrum. That characteristic is very valuable in the sense that workload requirements can change over time. For example, a bottleneck that exists today could change substantially over a period of months, transitioning from being processor-bound to memory-bound to I/O-bound. The flexibility of Intel architecture is extremely valuable in accommodating that variability.

Customers have consistently indicated that they want the same kinds of features and capabilities in micro servers that they have come to expect from traditional rack-and-blade server infrastructure. In particular, they need support for 64-bit software, virtualization support, error-correcting code (ECC) memory, a full range of power options, and broad software compatibility. The Intel Xeon processor E3 family ranging from 95W to 20W offers all those advantages for micro server platforms. In 2012, Intel plans to launch sub 10W server-optimized processors based on Intel® Atom™ microarchitecture that will offer the same benefits.

Growing Industry Momentum for Intel® Architecture-Based Micro Servers

Industry momentum is growing in the micro server segment among OEMs, with an array of new product offerings from industry leaders. Broad-based innovation includes micro servers based on the full range of platforms, mostly using Intel Xeon processors but also including some unique innovations based on today's Intel Atom processors. A few notable recent entrants to this market segment include the following:

- **Dell DCS5120**, incorporating 8 to 12 Intel Xeon processor-based nodes in a 3U chassis.
- **SeaMicro SM1000-64**, incorporating 512 Intel Atom processor-based nodes in a 10U chassis.
- **Tyan FM65-B5511**, incorporating 18 Intel Xeon processor-based nodes in a 4U rack-mount enclosure.

SMALLER ISN'T ALWAYS BETTER: A WORD OF CAUTION FROM GOOGLE*

Notwithstanding the value of micro servers, most workloads continue to favor a relatively lower number of more capable server systems. In his paper, "Brawny cores still beat wimpy cores, most of the time,"[1] Google Fellow Urs Hölzle warns that applications being designed for many smaller cores must be chosen carefully. Specifically, this model carries the risk that the workload could require more performance in the future but cannot be parallelized further.

Moreover, the software development complexity of massive parallelization may offset the cost savings in capital and operating expenses. Networking is another issue, since by deploying a larger number of less powerful servers, the number of ports required and switching overhead increase substantially, potentially impacting management and performance. The process of dividing workloads into many small pieces also carries the risk that one unit of work may leave the resources on a node under-utilized but unable to take on another unit of work, decreasing overall efficiency.

Standardization facilitates this breadth of innovation in micro servers. To that end, Intel worked with other companies, such as Quanta and Tyan, in the Server System Infrastructure (SSI) Forum to create the Micro Module Server specification,[2] which was released in January 2011. SSI, a trade group based in San Ramon, California, works with industry leaders to drive infrastructure standards. The forum's goal is to enable market growth through the standardization of computing technologies, emphasizing power efficiency and cost reduction for micro server end-customers.

The Micro Module Server Specification is a flexible node building block that is explicitly designed to be easily tailored to a wide range of power and performance design points. The specification defines the module dimensions and pinouts, but allows OEMs significant flexibility to customize system architecture including backplane, interconnect, and storage configurations. It uses broadly available, industry-proven, low-cost connectors and requires module mechanical interoperability while leaving electrical interoperability optional.

Overall, this approach strikes a balance between flexibility and standardization by facilitating product differentiation by OEMs within their product lines and relative to competitors while maximizing design reuse. As such, the specification fosters a number of OEM benefits, including the following:

- **Address varied product requirements** by innovating within a standards framework in areas such as system form factor, chassis, cooling, switching, and management.
- **Maximize ROI** by taking advantage of system design investments in multiple products that target different applications and market segments.
- **Accelerate time to market** by taking advantage of design efficiencies associated with the specification itself, existing design work, and Intel validation and design collateral.

End-customers are the eventual inheritors of this design flexibility and efficiency, which benefits them in terms of innovative product availability to meet evolving business needs with low cost and high ROI.

Conclusion

Micro servers based on Intel architecture offer a new option for the deployment of compute solutions that are particularly well suited to the needs of lightweight scale-out of non-enterprise applications. This topology can provide efficiency benefits by reducing system infrastructure such as fans and power supplies.

KEY MICRO SERVER REQUIREMENTS:

- Low power
- 64-bit processors
- Intel® Virtualization Technology support
- ECC memory
- Broad software compatibility

Micro servers may also enable customers to increase density, providing new ways to improve the use of existing rack and power infrastructure. As the industry begins to incorporate micro servers into existing and new usage models, they have the potential to improve the TCO associated with supporting this particular class of workloads.

[1] Hölzle, Urs, *Brawny cores still beat wimpy cores, most of the time.* http://static.googleusercontent.com/external_content/untrusted_dlcp/research.google.com/en/us/pubs/archive/36448.pdf (PDF)

[2] Yeraswork, Zewde, *SSI Forum Releases Micro Module Server Specification*, CRN (2011). http://www.crn.com/news/components-peripherals/229000520/ssi-forum-releases-micro-module-server-specification.htm;jsessionid=18NGOEkb0OTI5-cBnYSN7g**.ecappj01

0411/KH/MESH/PDF 325345-001US