



Motherboard-based Servers versus ATCA-based Servers

A comparison of costs, features and applicability for telecom application hosting

Summary:

After many years of struggling for market acceptance, it now appears that ATCA-based servers are beginning to gain traction in telecom applications. This is particularly true in larger applications with significant I/O capabilities that have traditionally been implemented via proprietary hardware or CompactPCI. As telecom equipment providers consider upgrading their older systems, the choice of standards-based ATCA offers many advantages.

At the other end of the complexity spectrum are motherboard-based servers that have become pervasive within both telecom and datacom applications. These typically 1U-tall rack mount servers are high-volume, cost-effective and versatile and are utilized for many applications. Motherboard-based servers are particularly well suited for applications that need minimal I/O capabilities other than Ethernet ports.

Both ATCA and motherboard-based systems utilize many of the same components, including Intel Xeon CPUs, DRAM memory sticks and common disk drives. Both types of systems are rack mounted, have similar rack densities and require the same amount of power. Operating systems and software applications run identically on both. Both have the same user interfaces and operate at the same performance. The remaining technical differences can be subtle, and have to do with the way the common elements are organized within the system.

Overriding the technical difference subtlety is the huge volume disparity between the two architectures. Whereas server-class motherboards are produced in the 10's of millions of units per year, ATCA CPU boards are produced in the 10's of thousands of units per year. This several orders of magnitude difference in volume – as well as some of the extra features required by ATCA – results in the price of ATCA-based systems being approximately double the price of equivalently configured motherboard-based systems.

So, when does it make sense to use ATCA-based systems versus motherboard systems? For primarily compute applications with modest I/O, motherboard-based servers offer the greatest cost-effectiveness. System configurations that require extensive amounts of I/O or I/O processing (i.e., many T1, T3, or OC48 ports, multiple DSP boards, etc.) may be good candidates for ATCA-based systems.

A hybrid system alternative – where motherboard-based servers are used for the CPU functionality and ATCA sub-systems are used to house the I/O boards – is also available. New PCIe over cable technology allows these hybrid systems to function at the same full performance as fully integrated ATCA systems. This hybrid system architecture can provide the best of both architectures: the cost-effectiveness of the motherboard server and the I/O capacity of the ATCA system.

Price comparisons of ATCA-based servers versus Motherboard-based servers:

Commodity-oriented motherboard-based servers are typically priced significantly lower than ATCA-based servers. OEM volume pricing for servers based on the two architectures is as follows:

1U motherboard-based server: \$2,900
ATCA 2U server with one CPU board: ~\$5,500
Delta cost: \$2,600 per unit

Since some telecom applications require redundant CPU functionality, it is also useful to consider the redundant pair configuration:

Two 1U motherboard servers: \$5,800
ATCA 2U server w/ 2 CPU boards: ~\$10,000
Delta cost: \$4,200 per pair of CPUs

Since most telecom OEMs provide thousands of server-based applications per year, the delta cost to utilize ATCA-based systems versus motherboard-based servers can be several million dollars per year.



One Stop Systems 1U NEBS server
(Based on Intel 5500 motherboard)



Configuration notes: both the motherboard- and ATCA-based servers include:

Two quad-core Intel 5500 (Westmere) CPUs
8GB DRAM
2TB hard disk drive
Dual 48V inputs
Full NEBS certifications

Pricing data is based on motherboard server pricing from One Stop Systems and pricing estimates for ATCA systems produced by Emerson Network Power. Quantity 1,000 pricing is assumed.

Similarities between ATCA-based and motherboard-based servers

There are many similarities between the two types of systems. This is because most of the important elements are identical.

CPUs: Typically both architectures support up to two Intel Xeon-class multi-core CPU sockets and support a wide range of available CPUs. ATCA blades are limited by the total power allowed per blade and due to the reduced height available for the CPU heat sink. Thus, the high-end of the CPUs is typically not available within ATCA. Motherboard systems provide approximately twice the power to the motherboard and have more room for better heat sinks, and thus can handle the full range of Intel CPUs.

Systems supporting the latest CPUs tend to become available 6-9 months earlier in the motherboard form factor versus ATCA. There are also more choices for motherboards than ATCA CPU blades.

DRAMs: Both use the same DRAM sticks (typically DDR3), and have the same number of DRAM sockets (typically 12) for a maximum capacity of 96GB DRAM.

Disk drives: In general, both architectures utilize standard (or server-grade) 2.5" and/or 3.5" disk drives and both support SSD drives. Many ATCA systems are constrained to only support 2.5" drives. This limits the disk drive capacity and performance within ATCA.

Operating systems: The wide array of operating systems that operate on PC-compatible servers will typically operate identically on both motherboard- and ATCA-based systems. Motherboards are somewhat better supported with drivers and more thoroughly tested with a broad range of OS's than ATCA due to their significantly greater market adoption.

Application software: Runs the same with the same performance on either architecture.

Performance: Same for both assuming same CPU, DRAM, disk, OS and application S/W.

Physical size/ mounting: Both are typically rack-mounted and both support a full array of rack mounting slides and other options.

NEBS certifications: Both types of systems are capable of full NEBS certification.

Ruggedness, environmental robustness: Both types of systems exceed all the telecom requirements. Motherboard servers are also used extensively in more rugged mobile and military applications.

Differences between ATCA-based and motherboard-based servers:

There are a number of differences between the two types of servers, including:

I/O board capacity: ATCA-based systems have the ability to contain a great deal of I/O capacity. This comes about from several features, including:

- A large blade size – nearly the size of a motherboard
- Ability to house up to 16 blades per chassis
- Ability to add AMC mezzanine boards to each card to add functionality without taking up another slot.
- Plenty of power and cooling for large boards with mezzanines installed

In comparison, typical motherboard-based servers contain relatively little I/O expansion. 1U systems have 1 or 2 expansion slots for the much smaller PCIe add-in boards. 3U motherboard-based systems contain 7 add-in board slots.

Rack densities: Motherboard systems are typically 1U tall. ATCA systems range in size from 2U to 13U tall, and can contain multiple blades per chassis. A 2U ATCA chassis contains 2 blade slots whereas a 13U ATCA chassis contains up to 16 slots. Thus, lightly loaded ATCA systems have slightly less rack densities, whereas fully loaded ATCA chassis can have slightly greater rack densities versus 1U motherboard-based servers.

In data centers, where rack density can be a stronger consideration, blade servers are commonly utilized. These blade servers are not based on ATCA for a number of technical reasons, particularly cost, power consumption and rack density. Data center blade servers can reach densities of 3X the density of 1U motherboard systems, and do so in a cost-effective manner.

Removable versus fixed CPU boards: The ATCA CPU blade is designed to be easily removable from its chassis, whereas a motherboard is more permanently fixed within its chassis. This characteristic provides both pros and cons for the two architectures. On the positive side, removable CPU blades are easier to replace in case of a failure. In a redundant system, it is possible to remove live boards from a system while it continues to operate. On the negative side, easy board removability encourages field personnel to remove and handle the boards. The removable ATCA blades are not enclosed in any protective casing, and thus when removed from the system are vulnerable to damage by handling (knocking off small caps or resistors or bending/breaking pins) and from ESD (simply touching ESD-sensitive components can damage the parts or significantly shorten their remaining operating lifespan). A vendor can attempt to minimize the handling and ESD risks through extensive training, but exposing the key elements of the system to unsupervised field personnel adds a system reliability risk. By contrast, motherboard-based systems are field replaceable at the system level. Thus, the systems are not intended to be disassembled in the field and do not present the same field quality risks.

Power inputs: For traditional Central Office telecom installations, both operate off of dual -48V inputs. Motherboard systems also have AC input versions. ATCA chassis require -48V feeds going into its backplane, and thus these systems are typically are not available with AC-input versions in the same physical size.

Central backplane – the ATCA systems Achilles heel: The central backplane within ATCA is the key element that allows multiple blades to be interconnected. This backplane is hard-mounted in the middle of the chassis, and all blades and power feeds interconnect through the backplane. The requirement of the central backplane provides a number of negative aspects for ATCA systems, including:

- **Single point of failure** – a backplane failure can take down an entire system (including multiple redundant CPU boards)
- **Bent pin failures** – fairly common with the type of backplane connector utilized in ATCA. This type of failure is particularly difficult since the backplane cannot be easily removed from an ATCA system.
- **Extra cost** – both for the backplane and the connectors required on both the backplane and the blades. The connectors utilized within ATCA are high-end specialty products. They cost typically 10X the cost of commodity PCIe connectors utilized on motherboards, yet provide the same performance characteristics.

Fault-tolerance through redundancy: Both system types offer fault-tolerance through redundancy. In the motherboard case there are simply two separate systems. In the ATCA case, there are typically two CPU blades contained in a single, shared chassis. Shared resources with the ATCA chassis include the internal power wiring, the backplane, the fans, etc. The possibility of failure within the shared resources of the ATCA chassis make this configuration somewhat less fault-tolerant (with many single points of failure) versus the dual-motherboard system approach. An easy way to visualize this difference is to consider what happens if the CPU blade shorts out. In the ATCA-case, both CPU boards in a chassis go out as the main power feed fuse is blown. In the motherboard server case, if a CPU shorts out that server goes down, but the other server (which is on a different power feed) continues to operate normally.

Market perceptions of ATCA-based versus motherboard-based servers

Does anyone really care?

Some do, but most service providers don't. Instead they focus on what their customers want, which is high-performance functionality at low prices.

With regard to telecom service providers, the high priority buying criteria include:

- Low cost

- High functionality
- High performance
- Long-term reliability

If you ask a telecom service provider, "Would you pay thousands of dollars more for equipment that provides the same functionality and performance with the same reliability, if that system were based on ATCA versus the more common motherboard architecture?", the answer will be a quick "no!".

In terms of actual market acceptance, motherboard systems are extremely pervasive within telecom and datacom centers. The usage of motherboard systems within telecom central offices is growing rapidly due to the system's cost-effectiveness and modularity. As telecom equipment becomes more and more dependent on easily deployed software compute platforms, motherboard-based systems will become more and more pervasive.

ATCA systems are currently relatively rare, but are expected to increase in popularity over the next few years. This increase will come primarily as ATCA replaces older proprietary and CompactPCI architectures. In its formative years, ATCA was heavily promoted as the next great technology for telecom equipment. As the reality of the marketplace has become evident, it is now clear that ATCA will play a niche role as one of several architectures that are useful for telecom equipment.

What causes such a large pricing gap between ATCA and motherboard-based systems?

There are a number of factors that cause ATCA systems to be significantly higher priced than equivalent motherboard-based systems, including:



Intel 5500 motherboard



Emerson Network Power ATCA CPU blade

Volume: As stated previously, motherboard systems are produced in the 10's of millions, whereas ATCA systems are produced in the 10's of thousands of units per year (sources: VDC

Research and Gartner Group). The key volume drivers for the motherboard volumes are in offices and data centers, where they are pervasive. Since ATCA is oriented toward telecom and is fundamentally expensive for compute-only applications, it cannot take advantage of the much larger volume marketplaces.

High volume has a profound effect of prices. And a two or three orders of magnitude difference in volumes creates a night-and-day difference between the two architectures.

With high volumes, manufacturing becomes much more efficient. The manufacturing process becomes highly automated, with every step hard-tooled for efficiency. Investments made to improve costs – value engineering, automated assembly, hard-tooled equipment, automatic test equipment, etc. are easily justified with high volumes. Further, the costs associated with the original board design as well as bios and firmware development can be amortized over a much larger volume.

Commodity versus non-commodity vendor businesses: High volume of a relatively non-differentiated product (Intel provides the motherboard reference design for free) creates a vendor base that is geared toward the production of a commodity product. Vendors of commodity products gear their entire business model toward cost efficiency.

	<u>Commodity (Motherboard)</u>	<u>Non-Commodity (ATCA)</u>
Supplier margin	10-15%	40-50%
Packaging	Optimized	Specialized
CPU and DRAM	Sold separately	Bundled
Sales channels	Broadly distributed	Limited distribution
Software license fees	Fully absorbed	Priced into product
Competitive alternatives	Extensive	Minimal

Extra components: The requirement for extra components within ATCA also adds extra costs. These extra components include the backplane, the extra connectors required on each blade, and the board removability features. In addition, ATCA requires a particularly complex and expensive system management microcontroller on each blade. The cost of this implementation is primarily in the firmware (often provided by a third party), which due to the modest volumes of ATCA, remains a significant per unit cost.

-48V orientation: ATCA was designed as a telecom Central Office-optimized architecture. The key feature of this optimization is that dual -48V feeds are required going into the backplane for power distribution to all the blades. This power requirement adds costs to each blade. Every blade in an ATCA system needs to contain circuitry to combine the two power input feeds

(ORing circuit), filter the power, and convert the voltages to useful levels. In installations where AC power is preferred (currently about 15% of telecom equipment installations) ATCA is particularly expensive. In these applications, a separate power supply (or for dual feeds, two external power supplies) is required to convert AC to -48V, so that it can power the ATCA chassis.

The best of both – Hybrid systems

Given the strong economic advantages of motherboard-based systems and the large I/O capacity of ATCA systems, some telecom equipment designers are building hybrid systems as a means to achieve the best of both worlds. In a hybrid system, one or two motherboard-based servers are used to front-end an ATCA system that contains only I/O boards. The connections between the motherboard systems and the ATCA system are a combination of PCIe and Ethernet. Both interfaces are connected via cables over to the ATCA I/O sub-system. In this configuration, a special ATCA backplane or interface board is required to bring the PCIe and Ethernet connections out to the rear of the ATCA enclosure.

Hybrid systems also allow the consideration of other architectures for the I/O sub-systems. With the CPU decoupled from the I/O, many optimizations are possible. In order to minimize costs, an I/O sub-system could be based on a single board design mounted in a low-cost 1U I/O chassis. Another alternative is to utilize a low-cost PCIe add-in card expansion chassis which can hold up to 20 standard add-in boards. This solution provides excellent modularity and scalability, while preserving the commodity advantages of standard PCIe add-in boards.



A hybrid system utilizing dual motherboard-based servers and an ATCA-based I/O sub-system

Summary Comparison:

	Motherboard-Based	ATCA-Based
Similar Features:		
CPU	Dual Intel Xeon CPUs	Dual Intel Xeon CPUs (Limited due to power and cooling constraints)
DRAM	96GB	96GB
Disk drives	2.5" and 3.5" drives (multiple drive capacities)	2.5" and 3.5" drives (Some systems limited to a single 2.5" drive, which provides less capacity and lower performance)
Operating Systems	Windows, Linux, Solaris, etc.	Windows, Linux, Solaris, etc.
Application software	All applications that run on the above OS's run on this platform	All applications that run on the above OS's run on this platform
Performance	Same performance with same CPU, DRAM and disk	Same performance with same CPU, DRAM and disk
Physical size/mounting	19" or 23" rack mounting	19" or 23" rack mounting
NEBS certifications	Yes	Yes
Ruggedness	Yes	Yes
Differences:		
I/O board capacity	2-7 slots for PCIe add-in cards	2-16 slots for ATCA add-in blades (ATCA blades are much larger than PCIe add-in cards)
Rack densities	1 server per 1U	½ to 1.2 blades per 1U
CPU board mounting	Fixed	Removable
ESD and field handling	No field handling of exposed electronics	Blade removability encourages field handling of exposed electronics
Power inputs	AC or dual -48V DC	Dual -48V DC only
Central backplane	No	Yes (Causes single point of failure, bent pin failures and extra cost)
Fault-tolerance	Dual Systems	Multiple Blades within a common chassis (shared elements become single points of failure)
Pricing:		
Single server pricing	\$2,900	\$5,500
Dual server pricing	\$5,800	\$10,000
Market volumes	Millions	Thousands
Supplier types	Commodity – low cost	Specialty – high margin