

Whitepaper: Deploying Network Services: Network Appliances vs. AdvancedTCA

Matthew Dharm
Lead System Architect, JumpGen Systems LLC

24 January 2011

Overview

As network traffic continues to increase, the desire of organizations to deploy ever more complex services on the network increases also. These services can be anything from data collection and analysis to managing streaming applications to specialized applications such as algorithmic trading, HFT, low latency messaging, DPI, traffic shaping, filtering, lawful intercept, and beyond. While many of these services can run on “commodity” server hardware, many require specialized equipment consisting of various combinations of general-purpose CPUs, NPUs, and switch fabrics.

Often, the requirement for specialized hardware leads system designers to automatically select a telecommunication-specific platform such as AdvancedTCA. While ATCA has several key benefits for telecommunication-centric applications, it is also burdened by several drawbacks which can make it less competitive for many deployment scenarios that do not require these specific features. In these scenarios, a 1U Network Appliance form-factor is often more appropriate.

In this paper, we will compare ATCA with a Network Appliance approach. We will examine several technical aspects of the two platforms as well as examine the cost differentials for both small and large deployments. We will demonstrate that for deployments that do not need ATCA-specific features (such as -48VDC power, extreme modularity, or complicated management), a Network Appliance can offer more features in less space with reduced cost, as compared to a similar ATCA system.

Technical Comparisons

ATCA boards and Network Appliances both offer similar technologies for handling large volumes of network traffic efficiently, including a variety of CPUs, NPUs, I/O capabilities, and switches. However, they are designed with very different requirements in mind and as such have very different technical strengths and weaknesses. We will examine a few of those here.

Latency and Bandwidth

Latency (the total time for a single bit to traverse a system) and bandwidth (the total number of bits per second to traverse a system) are two related (yet different) parameters which factor heavily into system design in both ATCA and Network Appliance architectures. Finding ways to decrease latency and increase bandwidth can significantly improve system performance without the need for a costly processor upgrade.

Latency and bandwidth are both, generally, functions of several different factors. While some of these factors are common between ATCA and Network Appliances (for example, the processing time of a packet inside an NPU), other factors such as switch topology are very different between the two types of systems.

In the case of both latency and bandwidth, ATCA systems suffer from being constrained by a board-and-backplane architecture, where the interconnections between boards are limited by the choice of the ATCA platform. For example, in a dual-star backplane, each node board only has two backplane network ports; that limits the total bandwidth any given board can attempt to process. If a node board required multiple on-board processors to handle 2 network ports of bandwidth, then the node boards would also require a local switch that would add significantly to the latency of each packet. This issue exists even in a full-mesh design, where even though each node board has multiple network ports, each node has a local switch adding to the latency.

A Network Appliance, on the other hand, does not suffer from this problem. Bandwidth between processing elements is a design choice made relative to the capacity of any given processor and potentially the software architecture of the end-product. Since all processors and the switch are on the same board, fewer switches are needed which leads to reduced latency and increased performance.

Expandability

One commonly perceived benefit of ATCA is the ability to vary the number of node cards in a system based upon the requirements of an individual deployment. Deployments which need more processing capability can have more node cards, while smaller deployments only get a few node cards. While this approach has some merits, it also has some significant drawbacks.

One such problem starts with the chassis. The system now must be designed either to use a large chassis in all deployments (which wastes chassis space and cost on smaller deployments), or to use multiple sizes of chassis (which increases the integration costs and interoperability issues). There are also questions of power and cooling in the chassis, especially the need to consider the change in airflow as various boards are added and removed. Similar considerations must also be made for bandwidth of the switch cards, available I/O ports, etc.

The Network Appliance, however, attacks this problem differently. In this model, the appliance is a self-contained system; deploying more capacity simply requires adding more appliances. Each appliance is well-balanced for power, cooling, bandwidth, and I/O capacity. Each appliance can be linked to other appliances in a deployment through a variety of technologies, ranging from 10G or 40G Ethernet to low-latency Infiniband or beyond.

Power and Cooling

ATCA is somewhat notorious for power and cooling limitations. The power limitations were significant enough that PICMG (the organization that manages the ATCA specification) actually increased the per-slot allowable power by 50% during a recent revision to the document.

A given ATCA board is now allowed to consume up to 300W of power. Unfortunately, as every system designer knows, consuming that much power is much easier than cooling that much consumption. In a typical ATCA system, a bank of fans (possibly in a push-pull configuration) distributes intake air among the slots in the chassis; any given fan generally pushes air to 3-5 slots at a time. Unfortunately, in a configuration such as this, the relative air resistance of each individual board relative to neighboring boards is important; a board with unusually high or low resistance will get either very little or almost all the air from a given fan. Neighboring boards will either get excessive cooling (starving the high-resistance board) or too little cooling; either way, one board or the other is doomed to overheat.

The primary cause of this problem is the physical size of the cooling unit (i.e. the fan) being a significant multiple of the physical size of the heat generation (the board) when the variability between boards is high. The Network Appliance, however, does not suffer from this problem. It is designed as an integrated unit, complete with thermal analysis to demonstrate proper performance up to (and possibly beyond) the required specification. Where an ATCA board has a power maximum of 300W, a properly designed 1U appliance can easily go much higher.

Product Life Cycle

Life cycle (the time from product introduction to end-of-life notification) is a key reason many projects select ATCA as a platform. ATCA boards, in general, have a 7-10 year life cycle; this makes deployments over time and replacing failed units much easier for systems that are expected to be deployed for a long life cycle. However, this strength is also a weakness, as this requirement limits part selection to only those parts which are guaranteed to be able to meet that life cycle requirement. This often means that many other parts, which might be technically more suitable to the task, are not available even for projects which do not require the long life cycle.

A Network Appliance design, on the other hand, gives the flexibility to decide, based on the end-project requirements, what the life cycle needs to be and to design according to that requirement. Deployments which are constantly “evergreen” (i.e. upgrading to the latest technology on a regular, short schedule) can therefore have access to a much wider selection of parts to select from to create the optimal design. Deployments requiring long life, on the other hand, can select suitable long-life parts; the system designer has the flexibility to determine if long-life is a requirement.

CPU and I/O Density

While ATCA offers hot-swap capability of individual components of a system, this feature comes at a significant cost. Each board is very limited in the total space available, and much of that space is dedicated to re-implementing similar circuits on each and every board. For example, every ATCA board has a -48VDC wide-range power input, which is most commonly converted to tightly regulated 12VDC before being used to power on-board components. Also,

every board has a similar IPMC for management and to support hot-swap. All of these repeated components greatly limit the space remaining for processors and I/O through the small front-panel.

A Network Appliance, however, does not suffer from this problem. The network appliance is 17in wide and 1U high (1.75in), with a varying depth depending on deployment. The entire front-panel and much of the rear can be used for I/O connectors. With a depth that can go to several times that of ATCA, the available area for processors is quite large.

For example, consider a system consisting of two 10GigE switches and four NPUs (with associated DRAM). In an ATCA system, this is likely a 4-slot solution (two switch cards and two dual-NPU node cards), which would likely be a 5U chassis once power entry modules, shelf managers, and cooling were included. That system might only have a handful of 10GigE ports. A similar Network Appliance would only be 1U and could have over 24 10GigE ports on a 1U panel.

Interoperability Issues

Interoperability is, perhaps, the most obvious area where a Network Appliance has an advantage over AdvancedTCA. A typical ATCA system consists of a chassis (with power modules and shelf manager) from one vendor, switch cards from another vendor, and node boards from a third vendor. While ATCA attempts to standardize everything about a system that would affect interoperability, there are always gaps. The very existence of organizations such as the SCOPE Alliance is evidence of the continuing problems of ATCA in the area of interoperability.

While ATCA is theoretically multi-vendor with high-interoperability, in practice interoperability of everything from management to backplane fabric to airflow resistance creates significant problems for a systems designer. For example, while some ShMM functionality is required by specification, there is a significant variation in features and quality of the ShMM. Pinpointing responsibility for a fault, either in the design or in deployed units, can be difficult or impossible with so many different vendors involved in a single system.

Of course, an appliance suffers from none of this, because the entire unit is sourced from a single vendor. That vendor designs, manufactures, and then warrants the units, eliminating confusing debates between vendors about which component is at fault when a problem arises.

Cost Comparisons

ATCA systems are somewhat notoriously expensive, when compared to equivalent processing capabilities delivered in a different form-factor. This is driven by a number of factors.

AdvancedTCA Overhead

The “overhead” of an ATCA system is significant. By “overhead,” we are referring to those components of the system which do not directly contribute to increasing the performance of the system, such as the ATCA Shelf Manager (ShMM). While the ShMM offers many management features which are required by the ATCA specification, these features are often unnecessary for many systems. Other components, such as intelligent fans and power

modules also fulfill requirements of ATCA but add additional (possibly unnecessary) complexity to a system design.

A Network Appliance can be designed to include only those features which are required for a given application and thus can achieve significant cost savings over ATCA.

Integration vs. Customization

Another key area where ATCA has a higher cost is in the area of integration. This is not specific to ATCA; any platform which utilizes modular components from multiple vendors carries an integration burden. This burden can be so significant that many projects elect to purchase all ATCA components from a single vendor; while this reduces the integration problems in theory, it also greatly limits the potential list of components to choose from.

A Network Appliance is inherently designed by a single vendor that is responsible for delivering a fully-operational design in which all components interoperate properly with all other components.

Conclusion

It is clear that, for applications that require the full feature set of ATCA, the ATCA platform has much to offer. Features like extreme modularity, -48VDC power input, and complicated management features are popular in some market segments where ATCA can offer an advantage. However, there many applications that do not require all of the telecom-specific features, and so these features turn into burdens. The Network Appliance platform does not suffer from any of these drawbacks and can, in turn, offer some more advanced features. As a final point of comparison, consider the following table which compares the two approaches for a typical system featuring 2 switches, 4 NPUs, and a CPU:

	AdvancedTCA System	Network Appliance
Components	2 Switch Cards 2 Node Cards 1 CPU Card Chassis w/ ShMM, etc.	1 Box
Physical Size	8U, 19" rack	1U, 19" rack
Available I/O Ports	10-15 10G ports	24-28 10G ports
Internal Bandwidth	2 fabric ports per node card, shared among 2 NPUs	4 fabric ports per NPU
Latency	At least 2 switches of latency	Single switch
Cost	\$10k / switch \$12k / node \$5k / CPU \$6k / chassis Total: \$55k	\$27k, all inclusive