The Outlook for AdvancedTCA Platform and Blades

Featured Analyst: Karen Liu

October 2003

Table of Contents

- Executive Summary
- Introduction
- Initial ATCA Adoption Driven by Development Cost Savings
- Alternative Scenarios for Cost Development Reduction
- AdvancedTCA Applies to Wide Range of Telecom applications
- ATCA Adoption Will Accelerate If Materials Costs Fall
- Platform Adoption Rate Gated by Design Windows
- Commercial Blade Forecast is Elastic
- Additional Potential Benefits
- Conclusion
- Appendix A: Background, Assumptions and Forecast Methodology
- Appendix B: Forecast Details

Reference: Information regarding the AdvancedTCA platform can be found at *http://www.picmg.org*.

This report was undertaken by RHK under the co- sponsorship of Intel® Corporation and PICMG®.

Executive Summary

A new hardware platform specification called AdvancedTCA® (Advanced Telecom Computing Architecture) is gaining momentum in both the component supplier and telecom system vendor community. The specification from PICMG® is designed to bring the benefits of standardized building blocks to the data plane as well as control and application planes of system hardware. Standardization inside system products represents a major change to system vendors' business models. This study examines the first two of three potential phases of adoption: phases based on system vendor development savings, material cost reduction and carrier efficiencies respectively. Early adoption is proceeding well: system vendors have shown a surprising willingness to adopt the platform with only development cost savings assured initially. These system vendors have deemed the platform technically suitable for a variety of central-office based network elements particularly in wireless access, wireless edge and new access edge. The burden is now on the building block suppliers to drive a second phase of adoption by delivering on the promise of materials cost savings compared to in-house board costs. The tipping point, where adoption is sufficient to lead to materials cost savings, and in turn accelerated adoption, is estimated in this study by two bracketing scenarios. Assuming success, the market forecast for third-party commercial building blocks—including server blades, non-server blades, and common equipment—grows from a negligible amount today to $3.7B in 2007.

Introduction

The telecom industry historically viewed standardization as relevant only to interfaces between network elements. Now system vendors actively seek improved business models as they are squeezed between increasing technology costs and decreasing time windows to provide application support to carrier customers. Commercial off-the-shelf (COTS) building blocks including chassis, backplanes and a variety of server, network processing unit (NPU) blades, digital signal processor (DSP) and disk storage blades offer an attractive way for system vendors to reduce the effort of designing and building their products. But a number of factors will need to converge over the next several years for the success of AdvancedTCA, also called ATCA™, or any standard platform to be assured. This study describes the drivers, barriers and likelihood of ATCA penetration into seven segments of telecom equipment.

RHK interviewed a dozen leading system vendors and four ecosystem suppliers for this study. The focus of the paper is on whether the profound business model change associated with widespread use of standardized hardware is likely to occur widely rather than on evaluating the relative chances of success of various standards. The first section of this paper discusses early signs of adoption driven by system vendors' need to offload hardware development. The second section presents a detailed forecast for further adoption over the critical period of the next five years. Due to the early stage of ATCA and the number of factors that could influence its success, two scenarios are used to drive forecasts. The critical difference between scenarios is the ability to achieve volume-driven price reductions and thereby pull further platform adoption. The forecast concentrates on network element types that have emerged from early adoption trends as most likely to use ATCA. The paper concludes with a brief discussion of additional factors that might lead to further adoption.

Initial ATCA Adoption Driven by Development Cost Savings

RHK interviews revealed system vendor attitudes have shifted in the last few years toward greater openness to third-party vendors including the ATCA model. System vendors cite the following as the top three benefits they hope for when evaluating ATCA: 1) materials cost savings so they can be price competitive, 2) time-to-market and 3) development cost savings.

The decisions of the early adopters today have been motivated primarily by development cost reduction. Interviewees are not yet convinced that materials cost savings will eventuate: a previous standard from PICMG, CompactPCI (cPCI), saw only limited adoption and hence little volume price reductions. While they do not have the same doubts about time-to-market benefits, suppliers relate that some system designers today must actually delay development schedules as ATCA components roll out.

Today's priority on reducing development cost is one key reason ATCA stands a much better chance of success than earlier efforts standardize telecom hardware. The degree of prioritization determines the likelihood of system vendors to adopt ATCA early. Some have made strategic decisions to stop hardware development on a division-wide basis: 3rd party standard or custom hardware is now the only choice. Smaller system vendors may find that an all-proprietary hardware solution is out of their reach. But few question the benefit: even large system vendors who could afford to maintain hardware development may prefer to shift their effort towards software, perceived as giving better return on investment.

System vendors are more varied in their assessment of the challenges that standard hardware poses. The depth of business model change varies with network element type and the company or division's strategic strength. They must believe they can continue to differentiate their products and successfully compete. Some organizations have already put their differentiation in software; others have been able to partition their hardware neatly so that proprietary hardware can be isolated to a few boards. Still others believe that their strength is in product architecture and technology, hence standard hardware is too restrictive. Even those that favor the approach anticipate that significant organizational changes and learning will be needed to support the new business model.

Adoption by Tier 1 System Vendors

Contrary to conventional wisdom, RHK interviews find interest in ATCA is not confined to Tier 2 players. As expected, market-dominating vendors are more concerned with defending market share than gaining share: these are reluctant to give up the barrier to entry that their proprietary hardware gives them. Tier 2 hardware system vendors hope that adoption of standardized hardware equalizes the playing field: by lowering barriers to mixing their products with the competitor's installed base, and by reducing the budget needed to launch new platforms. However, where two or three Tier 1 vendors vie for share or seek to extend dominance from one segment to others, they stand to gain greater benefits from common platforms compared to smaller companies:

- Common platform leverage across a larger number of products
- Greater awareness of supply chain issues—inventory, turn-around time
- Strategic interest in moving up to service and support rather than selling boxes

RHK is aware of 6 Tier 1 system vendors that have either decided to adopt ATCA for specific new product designs or are in detailed evaluation stage expected to lead to ATCA adoption, and has heard but not confirmed a number double that. (This may seem like a large fraction of Tier 1 companies but considering the number of products each vendor has, even outright adoption by a dozen companies would still represent a small fraction of telecom products.) Two companies spoken to have taken an informed position opposed to the use of ATCA for their particular needs.

Alternative Scenarios for Development Cost Reduction

RHK considered four possible ways that system vendors could choose to reduce development cost. There has been industry-wide movement from Status Quo toward Outsourced already that supports rather than competes with early signs of movement toward Standardized. Meanwhile the Emigrated scenario looms as a future possibility. If there are delays in ATCA adoption, the likelihood of the Emigrated scenario increases. In the meantime, the prospects for a move to standardized platforms are more favorable than in the past.

- **Status Quo**: System vendors continue to perform their traditional development role, with the industry consolidating as those who cannot sustain the resources to do so are forced to withdraw. Competitive strategy relies on either hardware-related intellectual property or leverage of installed base of proprietary products. Current market leaders in strong financial condition strengthen their position.
- **Outsourced:** System vendors outsource board as well as associated ASIC development to custom chip manufacturers or contract manufacturers. System vendor competitive advantages come from architectural skills, clever system partitioning and keeping up with (and choosing the right) best-of-breed chips of all sorts. System vendors still need superior hardware/firmware knowledge, at least enough to work closely with the manufacturing partner. Those with the best skills could manage to persuade those partners to give them favored status.
- **Standardized:** System vendors adopt standardized platforms for hardware and lower layer software. Competitive strategy mandates increased efforts on application software, customer support and network design services. System vendors who are best able to differentiate, especially on operational benefits to their customers, despite the restrictions of standard hardware have time-to-market advantages on new features. Ostensibly levels the playing field but in fact favors system vendors with strong network and operational understanding and leading software teams.
- **Emigrated:** The same number of development engineers but a fraction of the cost by moving to lower cost labor markets. Competitive strategy similar to Status Quo scenario

with additional advantage coming from best being able to ramp skill set in such geographies and retain key employees.

AdvancedTCA Applies to Wide Range of Telecom Applications

A successful common platform must balance between specifications broad enough to encompass enough network elements to generate economies of scale without being so broad as to fit none of them well enough to be adopted. It is by no means guaranteed that such a solution exists, and this remains a point of concern to system designers.

However, the interviews have found that early adopter system vendors have already evaluated the ATCA platform successfully for a variety of network applications including wireless access, wireless edge, wireline access, and others. Some attribute this success to unprecedented participation by a range of system companies in the defining committees. Others highlight the multi-threaded technical origins of the proposal. Transport designs influenced many of the definitions, for example in the use of rear transition modules. The PCI community contributed lessons learned from compactPCI's limitations. In contrast, some other standard hardware platforms come from a blade server heritage exclusively.

Although the base specification is already approved, work on extensions continues. Some of the technical limitations cited below are being addressed by work in progress. Two categories of network element where there are specific technical concerns cited by system vendors in the interviews are:

- Routers in both Edge and Core segments. Current board depth of ~20" needed to support the amount of silicon required for packet-processing blades exceeds the board depth of ~12" in front of the midplane for ATCA. ATCA mechanical dimensions are based on strict NEBS (network equipment building system) requirements of standard central office equipment but service providers have not held routers to these requirements.
- Wired Access equipment. Board spacing to support high port-density leads to boards as close as ½" compared to the 1" spacing for ATCA. Equipment vendors may desire access product designs to share attributes between central office and outdoor plant versions.

An additional area that is not a strong match, and implicitly excluded from consideration in the forecast, is:

- Non-central office equipment. The intent of the ATCA specification is to support blades in shelves in racks that comply with NEBS requirements. For pure servers (e.g. network management servers) that could reside in data centers, the NEBS orientation of ATCA does not violate any requirements but they are unnecessary.

ATCA Adoption Will Accelerate if Materials Costs Fall

Standardization moves such as ATCA are cooperative phenomena in which success breeds success: adoption itself encourages further adoption and volume benefits, driving still more adoption. The tipping point—where adoption is sufficient to lead to materials cost savings, and in turn accelerated adoption –is not known. It depends not on fundamentals but on the cleverness of system designers and their COTS suppliers to discover means to aggregate volumes. Just as contract manufacturers have been positioned and motivated to find commonalities of parts and processes, COTS suppliers could provide value both to themselves and their customers by discovering and exploiting commonalities in network element design.

This study postulates two self-consistent scenarios to bracket the tipping point and the size of market under both scenarios. The industry is currently in the Limited COTS state. It could cross into the Widespread COTS state if it reaches the tipping point. Once in the Widespread COTS state, it would be expected to remain there barring any new forces.

Table 1: Assumptions Defining Two Forecast Scenarios

Scenario Name	Limited COTS	Widespread COTS
Adoption drivers	Development cost driven only	Materials cost and development cost driven
Platform penetration	Adoption in products less sensitive to materials cost uncertainty: low volume products or high software content	Wider adoption includes penetration where materials costs must be minimized
Building block price	Moderate price declines with volume due to limited aggregation efficiencies in COTS supply	Steep price declines with volume due to good ability to leverage small number of COTS offerings across multiple board designs
COTS board usage	COTS for computing functions, storage and common equipment. Proprietary boards for other functions	COTS widely used for all functions, degree of usage depends primarily on volume and resulting price per blade type

Source: RHK

Drivers and Barriers to COTS boards

Even after system vendors have agreed to use a standardized hardware platform such as ATCA, there is a separate make-buy decision for each blade type within a system. The primary factors that affect this decision are included in the forecast model explicitly:

- Need for proprietary differentiation of a given blade type. For example, greater use of server blades is expected than packet-processing blades.
- Cost of the commercial solution compared to building in-house

Other factors are very important to system vendors but are not included in the model because their impact is not yet clear.

- Shorter lead times for build-to-order systems without incurring the costs of stocking components. However, system vendors could get similar support from contract manufacturers and without the need to maintain multiple vendors for different types of boards.
- Reduced effort to keep up with technology advances. System vendors hope the COTS vendor will take care of making sure the new blade is compatible with the existing ATCA system design. Conversely, system vendors are concerned they might have an increased configuration management burden if they move to COTS but the vendors provide inadequate lifetime support.
- System vendors perceive COTS boards as giving time-to–market in product introduction. However, this motivation does not assure continued COTS volumes if COTS is replaced by cheaper in-house manufacturing shortly after the system product begins its volume ramp.

The ability of COTS suppliers to aggregated volumes to lower cost and price will be critical to whether the industry will move from Limited to Widespread over the 2004 – 2007 timeframe. Volumes are limited initially and the aggregation is not uniform. One challenge is that the greatest beneficiaries are those products or system vendors with lower volumes. However, by the same token these are least attractive products and customers for the board suppliers. Many standardization efforts have encountered the problem of proliferating board types that dilute the interoperability and aggregation. Yet, consider the effect if all board suppliers flock to the same few obviously attractive board types: 1) dilution of the volume for supported boards and 2) preclusion of adoption by those network elements requiring unsupported boards. RHK sees encouraging early signs of coordination between board suppliers motivated primarily by time-to-market and development cost minimization. Continuation of such coordination will be critical to ensure the volume aggregation value proposition.

Note that system vendors do have the ability to aggregate volumes without resorting to standards-based hardware. Thus, COTS aggregation must do better than these alternatives to gain lasting cost advantage over in-house manufacturing.

- Internally defined cross-product platforms. Common today at the ASICS level, such coordination within large system houses is being extended to the board level. Commonality across multiple products not only gives system vendors volume savings but enable proprietary features that require interoperation between their own products, creating pull-through sales.
- Industry consolidation such that more segments become dominated by companies which alone encompass most of the unit volumes for a given network element or family
- Joint ventures or other industry cooperation between a group of system vendors to create a shared but proprietary cross-product platform

Table 2: Network Element Market Segments and Forecast Assumptions

Segment	Network Elements	Characteristics Assumed
Wireless Access	BTS/NodeB*, BSC/RNC, Transcoder	Early adoption seen for BSC/RNC but lower adoption for BTS/NodeB expected due to less benefit from common blade types. Generally less adoption for access than edge due to greater concerns about density and material cost. Some pull due to desire for common platform with wireless edge.
Wireless Edge	MSC, HLR, GGSN, SGSN/PDSN, Billing server, Multimedia server	Strongest initial segment: design window coincided with ATCA availability, and good technical match. Strong early adoption in GGSN, SGSN/PDSN for which even proprietary designs have been based on re-purposed platforms (e.g. ATM switches). Conflicted views on HLR. Market structure of several contending Tier 1 vendors and importance of software features favors ATCA.
Wireline Access	DSLAM, CMTS, MxU	Strong concerns about materials cost cause large elasticity between Limited COTS and Widespread COTS scenarios. Assuming market structure remains as is, dominant vendor is unlikely to favor any change of platform, limiting uptake to smaller vendors. Believe desired advances can continue on existing platforms for least 3 years.
Edge	Edge router, Multiservice Switch, Optical Edge Device	Dominant vendors unlikely to favor any change in platform. Current platforms have headroom, so no design window for at least 2 years.
New AccessEdge	Media gateway, Softswitch, Media server	Incumbent vendors have desire to migrate softswitch off legacy voice-switch with minimum development. Both contending and incumbent vendors already using cPCI and likely to prefer larger ATCA platform. Flexible modularity of softswitch architecture is good match for ATCA. However, not an active design area as market growth remains uncertain pending wider application of softswitches.
Core Transport	Core router, SONET/SDH ADM,, WDM	Current platforms recently designed with headroom. Design window deferred at least 5 years
Signalling	Signalling server, STP, SCP	Technically a good match to ATCA, but large part of market held by established vendors with no incentive to redesign for a currently uncertain market

Source: RHK Inc.
*Adoption lower than rest of segment

Platform Adoption Rate Gated by Design Windows

In telecom, new platforms are adopted infrequently and reluctantly when the existing platform can no longer be extended by the addition of new cards or software upgrades; typically new releases of products can continue on the same platform over 5 – 10 years. Hence the most important factor in determining ATCA adoption rates for the 5-year forecast period was the timing of platform design windows. Adoption rates in the longer-term are expected to depend more strongly on the match between the platform and the segment in terms of both technical and market factors.

The system level forecast focused on over 20 product types identified as mostly likely to see major new designs. For this study, they were divided into seven market segments with distinct characteristics. Definition of the segments, included network elements and factors affecting the adoption expectations are shown in Table 2. Further discussions of forecast methodology and assumption details can be found in Appendix A.

A combination of industry and RHK network element unit forecasts were used as inputs. RHK assigned ATCA adoption rates based on the attributes of each segment. Figure 1 shows the degree of penetration by segment expected in 2007 under each scenario. For each segment, the % shown is the unit forecast of ATCA-compliant systems divided by unit forecast for the total number of systems in segment. The remainder in each segment is assumed to be proprietary designs (i.e. alternative standards-based platforms were not considered). The total number of systems per segment includes only those network element types explicitly included.

Figure 1: % of forecasted network elements converted to ATCA by segment

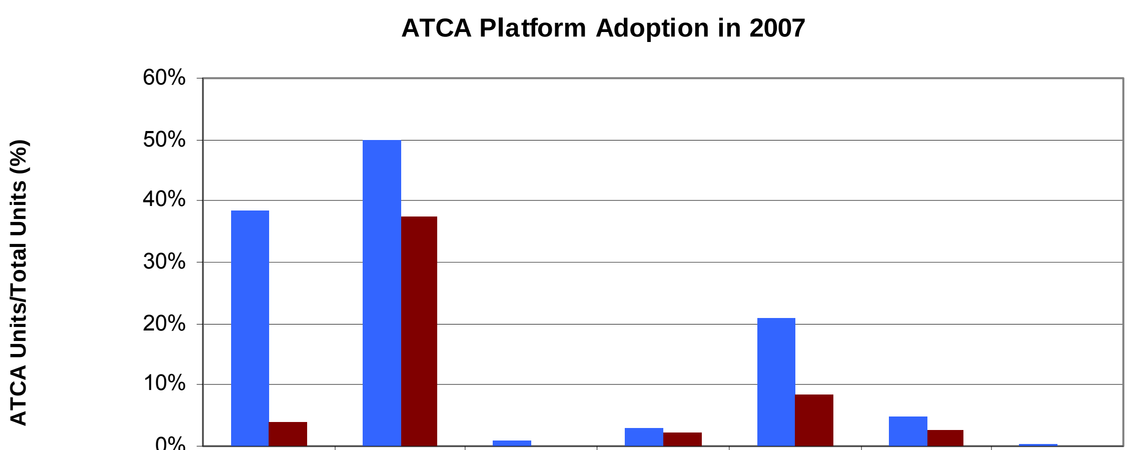

	Wireless Access	Wireless Edge	Wireline Access	Edge	New Access Edge	Signalling	Core
Widespread COTS	38.3%	50.0%	1.0%	3.0%	20.8%	5.0%	0.3%
Limited COTS	3.8%	37.5%	0.1%	2.3%	8.3%	2.5%	0.1%

Source: RHK

The network element forecast is believed to be conservative in a couple of ways. First, it starts with explicit inclusion of certain network elements rather than total capital equipment spending. As well, known but unproven potential upsides have not been included. These include additional equipment pull that would be driven in the event of widespread new services such as Push to Talk, Push to Connect, Push to Media, Interactive Media Services, Ring Back Tones, Video on Demand over Cable, Video on Demand on DSL, etc. Likewise, network changes such as substitution of softswitches for legacy local voice switches have not been assumed. Possible downsides to the penetration forecast include greater than anticipated slowdown in design activity due to financial pressures on system vendors.

Figure 2 Units forecast for ATCA network elements by segment

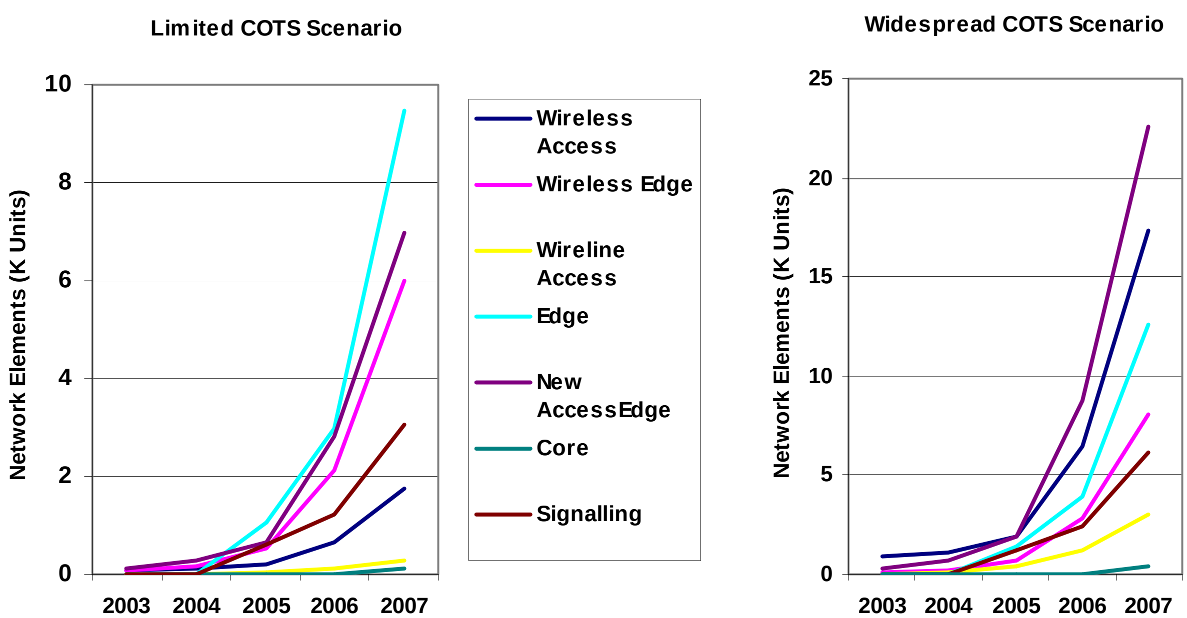

Source: RHK
Note: Values can be found in Appendix B, Table B-1

Figure 2 shows the absolute count of ATCA network elements for the purpose of showing the timing of adoption over the forecast period. Inferences about relative market size in revenue should be made with caution: network elements, even within a single segment, can vary widely in size from single shelf to multiple racks with a correspondingly wide range of price and blade composition. Similarly, the blade composition varies within each segment. The following section details the blade forecast derived from this network element forecast.

Commercial Blade Forecast Is Elastic

Both forecast scenarios assume that ATCA system vendors will use 100% COTS chasses, backplanes, server and storage blades because these correspond to items where either the number of such components per system is low, so the cost does not materially affect overall system cost and/or the system vendor recognizes there is no benefit in uniqueness, performance or cost to develop a proprietary version of such a blade. However, the Widespread COTS scenario assumes attractive pricing has driven greater platform adoption and hence higher overall volumes even for these elements.

Use of other blades such as NPU, DSP and advanced switch fabrics is more sensitive to scenario. Higher volume per blade type year by year is assumed to lead to price declines and COTS usage increases. The two scenarios differ in their assumptions regarding how steeply the price and usage vary with volume. The Limited COTS scenario exists today with cPCI and commercial (non-standard) processor blades. Early adoption of ATCA points to wider interest but still within the Limited scenario definition. Pricing today is reported to be comparable to in-house costs—not surprising given the early stage of the board market: system vendors can access evaluation versions now and hope for the first wave of commercial products late 2003.

Figure 3 Building block forecasts under Limited and Widespread COTS scenarios

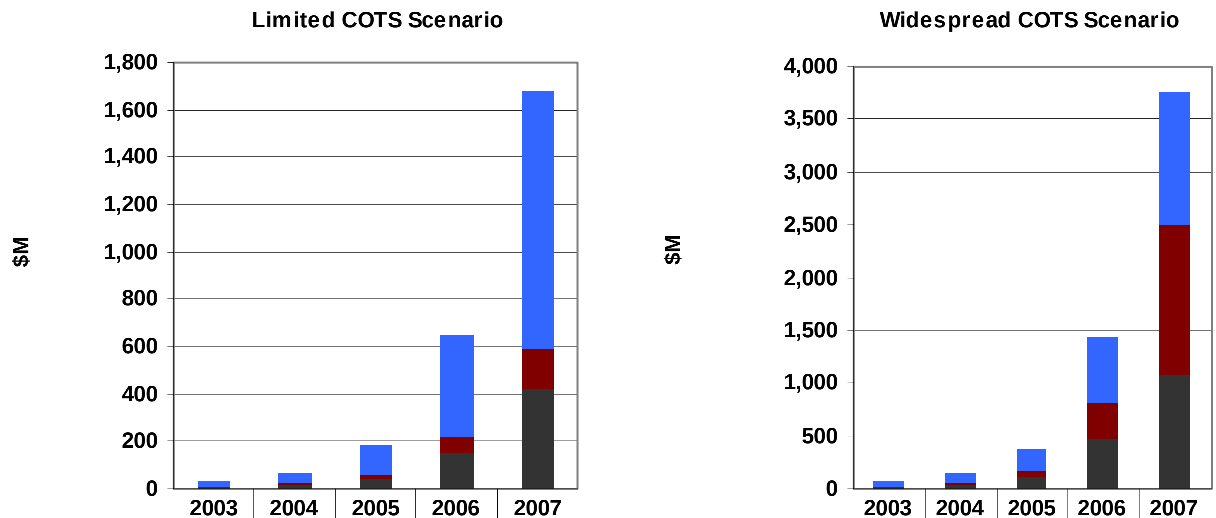

Limited COTS Scenario	2003	2004	2005	2006	2007
Commons	$28	$49	$121	$433	$1,085
Other Blades	$0	$6	$17	$65	$173
Server Blades	$7	$16	$46	$152	$423

Widespread COTS Scenario	2003	2004	2005	2006	2007
Commons	$50	$93	$208	$631	$1,266
Other blades	$1	$18	$67	$329	$1,410
Server blades	$19	$33	$105	$482	$1,086

Source: RHK

As adoption proceeds, both total volume and COTS usage % are expected to increase, leading to a steep increase in COTS board volumes (more pronounced for the Widespread COTS scenario). Both scenarios conservatively assumed the same critical price level of 20% below in-house alternatives for platform adoption and for maximum COTS blade usage. Maximum usage is assumed to saturate below 100% for boards likely to contain proprietary content such as DSP or packet processing NPUS. In reality, system vendors evaluate the combined cost amortized development and manufacturing against the price of COTS boards. However, for most successful products, the amortized development cost becomes a small fraction of the board cost. As well, system vendors express uncertainty about the residual cost to interact with the COTS vendors. The figure of 20% as a critical threshold to drive platform adoption comes from the system vendor interviews.

Note that the scenarios are not intended to represent the most pessimistic and optimistic cases. Nor is the state in 2007 the bounds of the market size. Given telecom product lifecycles of 5 – 10 years, ATCA will still be early in its lifecycle in 2007 with many design-in opportunities not yet reached. If the Widespread COTS level is achieved, ATCA will have reached its tipping point and will continue to grow strongly for many years after that.

Figure 4 Unit volume forecasts for server blades by segment for both scenarios

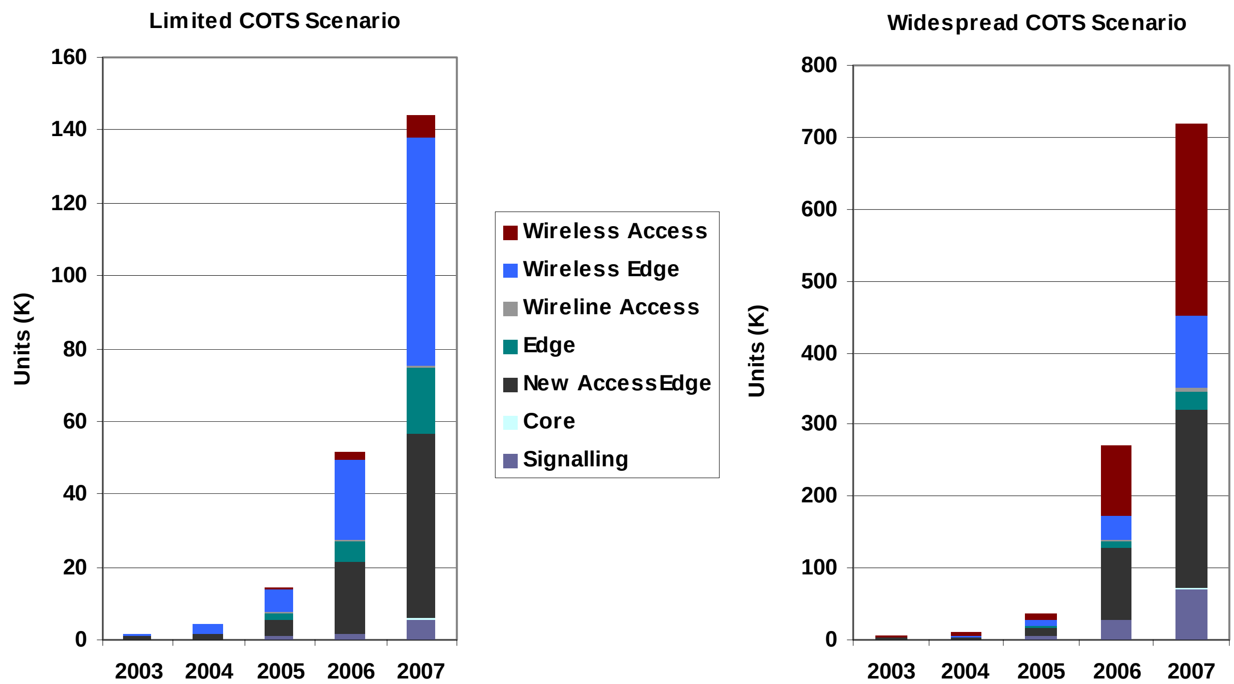

Source: RHK
Note: Assignment of blade types to server and non-server categories can be found in Appendix A, Table A-1
Note: Unit volumes can be found in Appendix B, Table B-2

The forecast model assumes the ATCA network element contains the same blades by functional type as the current proprietary systems plus formerly external associated hardware. The forecast model did not assume changes in density or architecture. Though such improvements could realistically be expected, it is difficult to project them in advance. One potential downside to the forecast is if greater than anticipated improvements in density reduce the number of blades per network element.

The model assumes ATCA network elements pull in disk storage and element management systems (EMS). System designers today use separate external hardware functions such as CO-located EMS and external disk storage simply because they believe the computer industry has superior expertise and volume efficiencies there. Their ideal would be the same attributes in a form-factor compatible with the rest of the system hardware. Other operational support systems (OSS) hardware that is not located in the central office was assumed to be unlikely to prefer ATCA compared to commercial non-NEBS servers.

The overall ATCA COTS component market including server blades, various other blades, and non-blade common equipment in the Limited COTS Scenario is over $500M by 2007. The non-server blades include NPU, DSP and disk storage blades. Non-blade common equipment is chassis, backplane, fans, etc. counted as a single unit per shelf. The majority of this market is in Server type blades. The forecast by year is shown in Fig 3.

In the Widespread Scenario shown on the left, while adoption starts with server blades, the DSP, NPU and Storage blades COTS usage undergoes a sharp rise as prices fall. By 2007, server blades are expected to make up somewhat less than half the overall ATCA COTS blade market. The total ATCA COTS blade market is around $2.4 B by 2007 in this Scenario. In the Limited COTS Scenario, the common equipment market is only slightly less than in the Widespread COTS scenario, but the blade market is just around $500M.

A COTS server blade vendor can expect to see the market develop by segment as depicted in Fig. 4. The market for other blade types is shown in Fig. 5. The detailed assumptions used to derive the COTS percentage of ATCA-compliant blades per scenario are discussed in Appendix A. A more detailed breakdown of blades into functional types is given in Appendix B.

Figure 5 Unit volume forecasts for non-server blades by segment

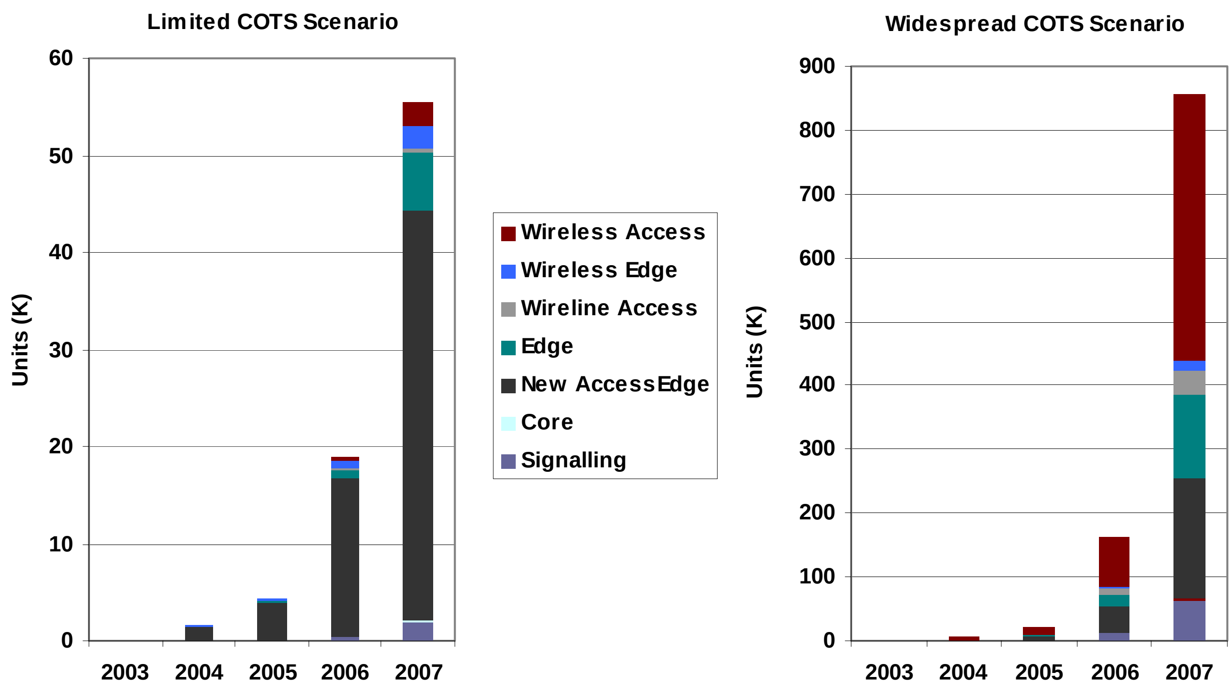

Source: RHK
Note: Assignment of blade types to server and non-server categories can be found in Appendix A, Table A-1
Note: Unit volumes can be found in Appendix B, Table B-3

Additional Potential Benefits

This study did not assume any improvements in density, innovation rate or other improvements that standardization might ultimately enable. The primary topic of this study is the interaction between system vendors and their suppliers. A complementary study would address the value proposition to service providers or carriers.

System vendors believe that standardized hardware gives greater benefits to their service provider customers than to themselves. Ultimately, if hardware becomes standardized, service providers could see seamless mixing of best-of-breed products from different system vendors with a uniform set of management and maintenance procedures. Yet, this nirvana for carriers might be hell for system vendors. System vendors fear the entry of companies with little telecom system design expertise leveraging standard platforms to quickly turn out adequate products, resulting in commoditization of the equipment market. If this happens, traditional system vendors might leverage their network and system expertise to hang onto the leading-edge market, but increasingly depend on solutions integration for their business. In a turn-around, they perhaps even gain returns on their intellectual property by becoming technology suppliers of leading-edge COTS boards.

As yet, there is little indication service providers will compel system vendors to move to standardized hardware. Both service providers and system vendors are still trying to quantify the benefits. In future, standardization might provide carriers with additional efficiencies as yet unproven:

- Ability to maintain best-of-breed network elements. Service providers evaluate and choose the optimum platform for initial deployment but are then locked into that vendor's subsequent releases. They have pushed network element interoperation efforts in response, but system vendor business models have precluded success. Perhaps standardization of underlying hardware will assist these efforts, but there is no indication yet that system vendors will not continue to differentiate on top of ATCA platforms.
- Ability to migrate network architectures. Changes at the network architecture layer can have greater impact than improvements in individual network elements. The timescale it takes service providers to adopt new network architecture is limited by lifetimes of the installed base of equipment. The possibility of re-purposing existing hardware, greatly increased by use of common ATCA components across multiple network element types, would shorten this timescale. But, it is unknown how re-purposing would work in the presence of continued system vendor differentiation. Likely, system vendors would need to be in control of the re-purposing, designing the hardware to be re-programmable only by them just as they plan for future releases on top of existing product today.

Conclusion

ATCA, a standards-based common hardware platform, has made a promising start. For continued success, the standards effort must balance the system vendors' need for differentiation with the COTS vendor need for commonality. The next several years will be a critical period. If COTS suppliers can meet and beat the price points required to trigger a second wave of adoption, it could be well on its way to changing the way the telecom equipment industry works. While such a profound change requires the convergence of a number of success factors and, most importantly, time, RHK believes the tipping point is attainable within the forecast period; if reached, the market can be expected to continue growing strongly for many years after that.

Appendix A: Background, Assumptions and Forecast Methodology

What is AdvancedTCA?

AdvancedTCA or ATCA is a hardware platform intended to specify pieces that meet requirements common to a large fraction of telecom network elements. Despite the word "Computing" in the acronym, this platform is designed to address the data plane as well as control and application planes. Telecom systems can be divided into two pieces: stuff common to all offerings within a segment, and truly differentiating hardware or software. Developing and manufacturing the common parts gives no competitive advantage to the system vendor: it can be considered "table stakes" cost that should be minimized.

An equipment manufacturer can choose to adopt ATCA definitions to a lesser or greater degree. A minimum level might be adopting the chassis and connectors and passive backplane but using proprietary content on cards and backplane switches. At the other extreme, the entire ATCA shelf complete with operating system and low-level software might be purchased as standard off-the-shelf. The base specification (http://www.picmg.org/specifications.stm) in PICMG 3.0 defines board form factors, connectors, backplanes, sub-racks or shelves. 48 V power distribution, NEBS 3 compliant mechanicals including cooling are also defined. Hardware management is also defined, including power-on sequencing, detection of faults, hot-swap and card-type detection. Management is a priority for the specification committee since the platform's credibility as carrier-class will be critical to its acceptance.

A growing list of PICMG 3.x specifications are intended to allow the equipment vendor a choice between commonly used flavors of switched backplanes capable of supporting up to 10G per card. The range of backplanes is intended to increase the types of network elements that can be addressed by ATCA: separate connections for control and data-plane.

An "ecosystem" of third-party building block suppliers (also called COTS vendors) offers to offload board hardware and firmware development from the system manufacturer as well. These board types include not only computing blades with general-purpose processors but also telecom-oriented blades with network processors, digital signal processors and a variety of interfaces.

Furthermore, although not required under ATCA definition, the broader open-standards vision includes similar efforts at higher layers. Those most likely to be relevant to ATCA systems include open-standard operating systems such as carrier-grade Linux and middle-ware components such as service availability APIs defined by the Service Availability Forum.

Development Option Scenario Discussion

Forces against Status Quo

Although strong incumbent system vendors, particularly those with dominant market share are motivated to maintain the Status Quo, technical and market forces will continue to increase against all proprietary hardware platforms. Advances in chip technology outpace system advances. Speeding up system advances to match is not economical if it implies replacing installed base of equipment. At the chip level, custom chip (ASIC) costs have increased dramatically as linewidths decline. 90 nm Si mask costs are typically $1 – 1.5 M, implying that only the highest volume telecom infrastructure products can now justify dedicated ASIC runs. New chip alternatives relying on a customization off the wafer such NPU, shared ASIC runs, or semi-custom ASSP allow proprietary hardware design to continue but point toward the Outsourced scenario.

With their carrier customers beginning to realize that they sell services not pipes, system vendors increasingly sell on OpEx savings and new service enablement, rather than box capacity. Hardware differentiation for enabling unique management and application software features may not require entire system to be proprietary.

Move to Outsourcing has occurred

Outsourcing of more and more design is consistent with an overall drift upwards in the food chain for both system vendors and their existing suppliers. ASIC suppliers already need to provide increasing design support in the form of building blocks and engineering staff to compete with off-the-shelf chips. The level of support for board level design could become increasingly important in how both ASIC and off-the-shelf chip vendors compete with each other. Likewise, for contract manufacturers to provide best manufacturing cost reduction to their customers, they have to influence the designs to best match their processes. While the design task has not been completely outsourced to date, the industry can be considered in a state favoring Outsourcing in general and the migration of more and more tasks to outsourcing partners can be viewed as a movement deeper into that scenario rather than a shift to a different scenario.

Standardized Likely to be Next

Over the next few years, the industry is poised to shift from Outsourced to Standardized. A move to outsourcing sets the stage for standardization. Like the COTS community, contract manufacturers are naturally in a position to discover commonalities between multiple products and multiple system vendors. Their business model is exactly to exploit found efficiencies, splitting the cost savings between their customers and their own profit. However, they are restricted by confidentiality agreements with each customer. An external standard provides neutral ground upon which to base such recommendations. Thus the contract manufacturers and the ecosystem supplying the common building blocks as off-the-shelf standard products are symbiotic.

Emigrated Scenario as either threat or accelerator to Standardized

However, the Emigrated scenario could also follow the Outsourced scenario. As time passes and adoption of ATCA is gated by system design windows and component vendor roadmaps, the likelihood of the Emigrated scenario increases.

Contract manufacturers already exploit their aggregated volumes to move manufacturing to low-cost geographies. At the design level, software development has moved to Asia with just a few companies (including reportedly 3Com, Cisco and Ciena) starting to look at hardware as well. As advanced chip fabrication moves to Asia, local chip design activity has already increased recently. Although not yet seen broadly, multiple threads could lead to board level design moving to Asia:

- Organic shift as Asian system vendors gain market share
- Global system vendors gradually build up design capability in-house in Asia

At this point, system vendors in North America and Europe are generally too concerned about loss of expertise to move their key development to Asia. However, less critical pieces—the same ones most likely to move to commercial standard product—may be eventually be moved or outsourced. Outsourcing to CM (contract manufacturer) or ODM (original design manufacturers) in Asia effectively combines the advantages of both Outsourced and Emigrated. This alternative is viable if system vendors really want to resist standardization. But, meanwhile, Asian system vendors are as interested as any in the economies offered by standard hardware. Should Asian vendors adopt standard hardware, others may be compelled to follow to match doubly reduced costs.

Platform Adoption Forecast Assumptions

The attractiveness of the ATCA platform, which varies by segment in terms of both technical and market factors, determines an asymptotic level of adoption (not reached within the forecast period). The timing also varies by segment. For the first five years, covered in this forecasts, the effect of timing is dominant.

Design Window Timing Model

While each TEM has its own product lifecycles and hence varying design windows, there do exist industry-wide bursts of design activity by segment. For simplicity, the model assigned a single nominal year for which product design using a new platform could occur. The assignment was based on the system vendor interviews as well as RHK's general industry knowledge. In some cases, different years were assigned for incumbent market leaders and for market contenders.

Industry-wide design activities are triggered by new standards such as 3G wireless or DOCSIS 2.0 in the cable modem area, or technology advances such as a broadly based increase in interface bit rate. For example, ATCA adoption activity to date has been greatest in wireless, believed to be due to high levels of design activity for 3G products.

Subsequent to the new platform design year, each segment was assumed to follow the same ramp in % on the new platform for simplicity.

Technical Attractiveness

Technical factors taken into consideration in the forecast model included:

- Similarity between ATCA and current implementations in terms of size, power, port density, etc.
- Maturity of product function. Well-defined products whose function varies little between competing system vendors are more likely to be well supported by a commercial ecosystem.
- Variety of blade types required. Needing a variety of blade types increases the attractiveness of buying less critical ones off the shelf. However, need for blade types unlikely to supported by the ecosystem weakens the attractiveness of the platform.

Little impact was found from maturity of product function. RHK believes this factor is more relevant for the commercial chip area than for boards-level platforms. In the New Access Edge segment (media gateway, softswitches) which is characterized by great product variation, the variability required does not seem to extend below the board level—in fact the blade platforms are well suited to supporting modularity and flexible blending of serve and data processing boards into a variety of system configurations. Several vendors in this area have already adopted cPCI for their softswitch products.

Commercial Attractiveness

Segments most likely to adopt the ATCA platform include those who for which lower volume system shipments, such as Edge rather than Access, lead to greater development cost pressure and somewhat greater tolerance to possible materials cost increases. Products who value in more in software than in hardware are also considered a better match because of good anticipated return on development cost re-allocated to software features.

CompactPCI, the 2.x family of PICMG specification, experienced telecom adoption limited by technical performance to smaller network elements, particularly with lower power dissipation per board. Adoption of cPCI was considered supporting evidence of ATCA suitability rather than explicitly included in the model.Even if the specification is sufficient for a given system product, a system vendor is likely to prefer a single platform that can span all versions of its product. Simply by virtue of larger board area and higher power consumption limits ATCA is more attractive and is expected to replace cPCI in many of its current applications.

Commercial Blade Forecast Assumptions

Table A-1: Functional Blades

Blade Type	Function	Typical Implementation
Element Master Blade (Server Blade Type)	Internal management of NE: control, alarm, processor, heartbeat functions; does not scale with traffic. Communicates with OSS	General purpose microprocessor
Services Processing Blade (Server Blade Type)	Heart of application server type system; scales with traffic or tasks; Co-located EMS blade in non-service NE	Powerful microprocessor, multiprocessor cluster or L7 NPU; memory, perhaps disk storage
Control-plane Blade (Server Blade Type)	Real-time processing of “slow-path” traffic and signalling; scales with traffic	Microprocessor, real-time OS, fast memory
Protocol Processing Linecard (Non-Server Blade Type)	Data-plane real-time processing of individual traffic streams based on specific protocols; aggregation; scale with traffic	Microprocessor or L2, L3 Network processor (NPU), DSP, Application Specific Standard Product chips (ASSP) by protocol; physical line interfaces
Signal Processing Blade (Non-Server Blade Type)	Various data-plane functions: radio domain filtering, speech transcoding, etc.	DSP
Packet Processing Linecard (Non-Server Blade Type)	Data plane real-time processing of individual packets including deep packet inspection, packet modification. Distinguished from protocol processing by less emphasis on external standards, more performance	L3 NPU or ASIC; physical line interfaces.
Interface Linecards (Non-Server Blade Type)	Physical layer interfaces only, no protocol or packet processing	physical line interfaces
High Performance Switch Fabric Blade (Non-Server Blade Type)	Directs data plane traffic control over backplane. More standard Switch Fabrics are assumed to be counted with Backplanes, and not located on distinct blades	High performance Switch Fabric chips
Storage Blade (Non-Server Blade Type)	Local data beyond memory capacity; ATCA in-system version of external disk farms	Disks

Source: RHK Inc.

Network elements were modeled as consisting of the following functional blocks, assumed for simplicity to be physically packaged in a separable blade. The forecast model is based on a functional breakdown (e.g. protocol processing) then mapped to implementation (e.g. microprocessor). The mapping from function to implementation varies widely and is expected to be fluid. For example, data-path functions may require different technology depending on bit-rate, e.g. rates below 1 Gbps can be implemented with general-purpose processors while higher data rates require DSP or ASIC. But as technology advances, these boundaries are expected to shift continually whereas the functional breakdowns remain more constant.

The COTS business model relies on being able to support the function of a large number of blade types from a limited set of individual products. This is possible by exploiting the combinatorial leverage of up to four mezzanine boards per blade. For simplicity, this forecast does not look below the functional blade level. As well, potential changes in system composition due to ATCA adoption— e.g. improvements in density—have been ignored.

In another simplifying assumptions, each network element type is assigned only a single typical implementation.

An element management system (EMS) converted to ATCA was assumed to consist of one additional Element master blade and one additional Storage blade, with a ratio of 1 EMS handling 30 network elements. Storage requirements were estimated per individual network element requirements with generally only 1 blade needed per network element except for video and multimedia servers.

Cost Model Assumptions

A generic cost model was assumed to cover a variety of network elements while acknowledging that real implementations differ from this model by up to a factor of 2x. These values are assumed stable over the forecast period with natural price declines offset by increasing function over time.

Table A-2: Cost Model

	Nominal Cost	Comment
List Price of Network Element	$100K per shelf	
Actual Sale Price	$60K per shelf	
Target Cost of Goods Sold	$35 K	40% gross margin, some NE actually up to 60%
Budget Cost of Shelf Pre- Integration	~$30 K	15% for integration
Cost Target per Board	$2 K	Divided shelf cost by 15: 12 – 14 blades + chassis & backplane

Source: RHK Inc.

Today, proprietary but third-party microprocessor-based computing blades are offered in the $3K to $5K range. System vendors have a mixed reaction to this price level. Many are willing to accept this premium for blades not within their core competency (e.g server or storage blades) but not for those that they either do well or that constitute a large fraction of the box (line cards). Others, particularly for segments with high price pressure are reluctant even for computing blades. The model assumes that all blade types follow this pricing except packet-processing and switch fabrics which are assumed to be 2x higher due to expensive chip content. The target price for the chassis, backplane and other common equipment taken together is assumed to be 10% of the total shelf cost.

Volume is clearly the required trigger as cPCI blades with less functionality and a smaller form-factor are still in the $2K range today after several years of availability. The model assumes a price decline of 10% and 20% per 5-fold increase in board volumes over 1000 units for the limited and widespread COTS Scenarios respectively.

The model assumes a doubling in penetration of COTS per 5-fold increase in board volume driven by associated price declines, reaching 100% at the critical price level of 20% below in-house alternatives. The methodology of the calculation starts with the low volume penetration rate. If the resulting volume is high enough to warrant the next level of penetration, the rate for that and any subsequent year was ratcheted up. No time lag between reaching volume and change in price or penetration was assumed.

RHK's model estimates that DSP and protocol processing boards can reach 80% penetration while packet processing will lag at 60%. All other blades types can reach 100%.

Table A-3: Price and Penetration Model

	Volume	<1 K blades	1K – 5K	5 – 25K	25 – 100 K	>100K
Limited COTS Scenario	Price	100%	90%	80%	73%	65%
	COTS/Total Boards	1.5%	3%	6%	13%	25%
Widespread Scenario	Price	100%	80%	64%	50%[1]	40%[2]
	COTS/Total Boards	6%	13%	25%	50%	100%

Source: RHK Inc.
Note: [1] price parity with in-house cost ,[2] threshold price of 20% below in-house cost

A simplified version of actual vendor roadmaps suggests the following timetable for when COTS products availability will cease to be a consideration limiting platform adoption. The model also assumes no COTS usage prior to these dates. Note that products may be available at earlier or later dates from vendors. For example, system vendors have little doubts that satisfactory server blades will be available and hence this issue is never an obstacle. Other blade types take into account an initial period of refinement before system vendors become comfortable with the maturity of those types of COTS products.

- In 2003: Microprocessor-based server blades
- In 2004: DSP, NPU and disk storage blades
- In 2005: Higher performance NPU for packet processing, high performance switch fabrics.

Appendix B: Forecast Details

Table B-1: ATCA network elements by segment (thousands of units) as in Fig. 2

Limited COTS	2003	2004	2005	2006	2007
Wireless Access	0.1	0.1	0.2	0.6	1.7
Wireless Edge	0.1	0.2	0.5	2.1	6.0
Wireline Access	0.0	0.0	0.0	0.1	0.3
Edge	0.0	0.0	1.1	3.0	9.5
New AccessEdge	0.1	0.3	0.7	2.8	7.0
Signalling	0.0	0.0	0.6	1.2	3.1
Core	0.0	0.0	0.0	0.0	0.1
Total	**0.3**	**0.6**	**3.1**	**9.9**	**27.6**
Widespread COTS	**2003**	**2004**	**2005**	**2006**	**2007**
Wireless Access	0.9	1.1	1.9	6.5	17.4
Wireless Edge	0.1	0.2	0.7	2.9	8.0
Wireline Access	0.0	0.1	0.4	1.2	3.0
Edge	0.0	0.0	1.4	4.0	12.6
New AccessEdge	0.3	0.7	1.9	8.8	22.6
Signalling	0.0	0.0	0.0	0.0	0.4
Core	0.0	0.0	1.2	2.4	6.1
Total	**1.4**	**2.1**	**7.6**	**25.7**	**70.2**

Source: RHK

Table B-2: Unit volume forecasts for server blades by segment as in Fig. 4

Limited COTS	2003	2004	2005	2006	2007
Wireless Access	0.2	0.3	0.5	2.0	6.0
Wireless Edge	0.7	2.4	6.1	22.4	62.9
Wireline Access	0.0	0.0	0.1	0.2	0.6
Edge	0.0	0.0	1.9	5.5	17.9
New AccessEdge	0.9	1.8	4.6	19.7	50.9
Signalling	0.0	0.0	0.0	0.0	0.2
Core	0.0	0.0	0.8	1.9	5.7
Total	**1.9**	**4.5**	**14.1**	**51.7**	**144.2**
Widespread COTS	**2003**	**2004**	**2005**	**2006**	**2007**
Wireless Access	2.3	3.6	9.8	98.2	269.5
Wireless Edge	1.0	3.2	8.5	35.3	98.6
Wireline Access	0.0	0.1	0.8	2.3	6.1
Edge	0.0	0.0	2.7	8.5	26.7
New AccessEdge	1.4	3.4	11.0	99.0	248.5
Signalling	0.0	0.0	0.0	0.0	0.8
Core	0.0	0.0	4.6	28.1	70.4
Total	**4.7**	**10.3**	**37.5**	**271.5**	**720.5**

Source: RHK

Table B-3: Unit volume forecasts for non-server blades by segment as in Fig. 5

Limited COTS	2003	2004	2005	2006	2007
Wireless Access	0.0	0.0	0.1	0.5	2.4
Wireless Edge	0.0	0.1	0.2	0.8	2.4
Wireline Access	0.0	0.0	0.1	0.1	0.3
Edge	0.0	0.0	0.2	1.0	6.0
New AccessEdge	0.0	1.5	3.8	16.2	42.4
Signalling	0.0	0.0	0.0	0.0	0.1
Core	0.0	0.0	0.1	0.4	1.9
Total	**0.0**	**1.7**	**4.4**	**19.0**	**55.4**

Widespread COTS	2003	2004	2005	2006	2007
Wireless Access	0.2	3.6	11.3	78.2	416.4
Wireless Edge	0.0	0.1	0.5	3.2	16.3
Wireline Access	0.1	0.2	1.7	8.2	35.7
Edge	0.0	0.0	2.0	20.3	131.8
New AccessEdge	0.0	1.2	4.9	40.4	188.0
Signalling	0.0	0.0	0.0	0.0	4.6
Core	0.0	0.0	1.5	12.4	61.9
Total	**0.2**	**5.1**	**21.8**	**162.8**	**854.7**

Source: RHK

Figure B-1 Detailed COTS blade breakdown by function

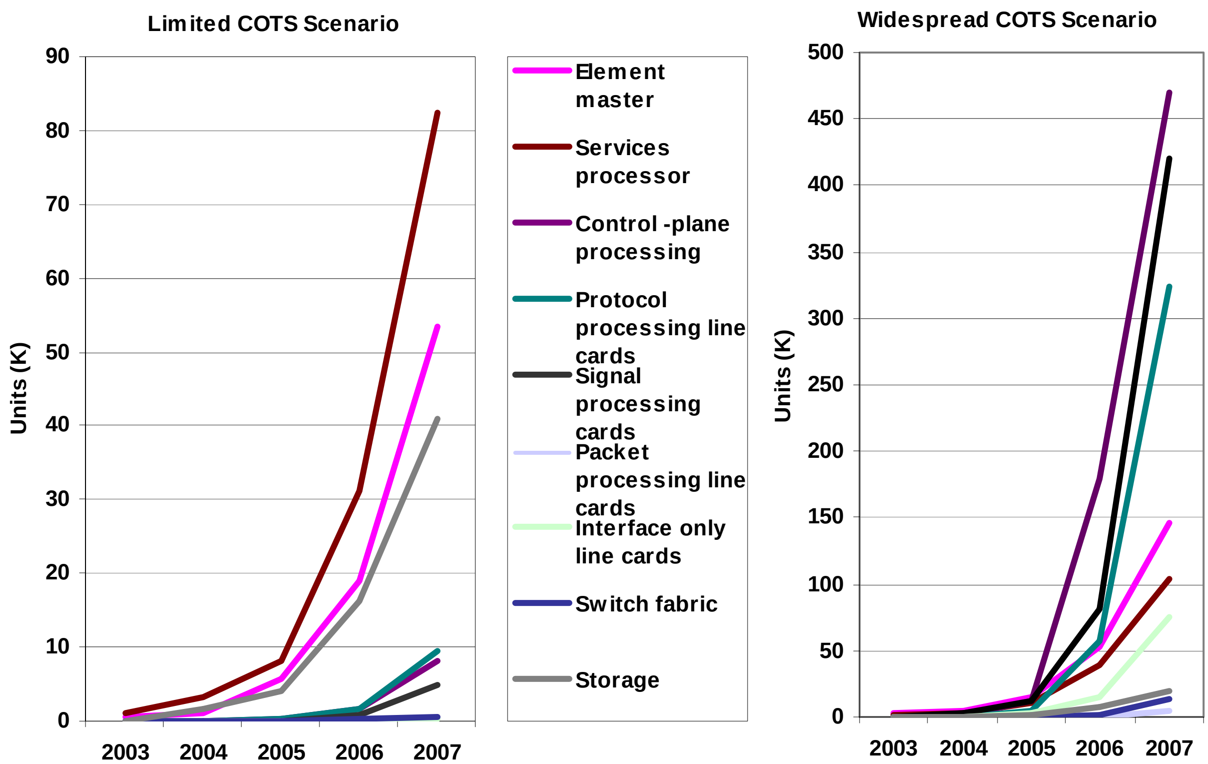

Source: RHK

Table B-4 Blade and Commons by Functional Type (Thousands Units) as in Figure B-1

Limited COTS	2003	2004	2005	2006	2007
Element master	0.6	1.2	5.6	19.0	53.5
Services processor	1.2	3.2	8.2	31.1	82.5
Control -plane processing	0.0	0.1	0.2	1.6	8.2
Protocol processing line cards	0.0	0.0	0.3	1.6	9.6
Signal processing cards	0.0	0.1	0.1	0.9	4.9
Packet processing line cards	0.0	0.0	0.0	0.1	0.2
Interface only line cards	0.0	0.0	0.0	0.1	0.2
Switch fabric	0.0	0.0	0.1	0.2	0.6

Limited COTS	2003	2004	2005	2006	2007
Storage	0.0	1.6	4.0	16.4	40.8
Shelves/ Backplanes	5.1	10.1	27.7	110.0	275.8
WIdespread COTS	**2003**	**2004**	**2005**	**2006**	**2007**
Element master	2.9	4.6	14.9	53.3	146.2
Services processor	1.3	3.8	10.0	38.5	104.3
Control -plane processing	0.5	1.9	12.6	179.7	470.0
Protocol processing line cards	0.0	0.4	5.0	57.4	323.3
Signal processing cards	0.0	3.7	12.2	81.2	420.4
Packet processing line cards	0.0	0.0	0.2	0.5	5.0
Interface only line cards	0.2	0.3	2.4	14.8	75.5
Switch fabric	0.0	0.0	0.4	2.1	13.7
Storage	0.0	0.7	2.1	7.7	19.1
Shelves/Backplanes	13.0	24.1	67.7	256.7	643.9

Source: RHK

