

ATCA for Telecom:  
Scalable, Flexible, Redundant  
& Cost-Effective Solutions

***Foundations for the All-IP Network  
Access Concentration***



[www.adax.com](http://www.adax.com)

WHITE PAPER

April 2011

## I. Executive Summary

Change and evolution in the world telecom network is a constant and exciting process. Wireless network technology has made steady progress since the days of AMPS and the 'brick-phone'. In the past 15 years we have deployed 2G, 2.5G, 3G, 3G+, and IMS networks. Now we are moving towards 4G, LTE, LTE-Advanced... into the future. Whilst it may be an old saw, it still has teeth: If change is constant then the only constant is change. In telecom we know this adage by heart.

The challenge to change is that development and deployment of new technologies and services is expensive. But change we must, even in these difficult economic times where the pressure to reduce costs conflicts with the necessary expense of innovation. Fortunately today we have options to choose from. Rack Mount Servers tuned for database and compute intensive applications are becoming cheaper and more powerful all the time. Leading the way for multi-purpose applications is ATCA, which provides a **scalable, flexible, redundant and cost-effective** platform architecture as depicted in Figure 1.

ATCA was specifically designed to accommodate wireless network growth. New markets require the most advanced technologies while maintaining legacy connectivity. Mature markets require upgrades to retain current customers and attract new ones, all the while keeping costs low by utilizing existing plant to its fullest. ATCA offers advantages to both these requirements; innovation and ROI. This White Paper will show the common needs of access technologies and use the Femtocell Gateway incorporating Access Concentration as an example of ATCA platform benefits. The node will be referred to as the FemtoGW/AC throughout.

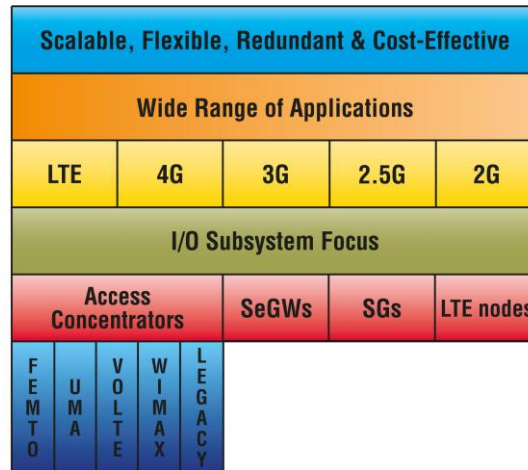


Figure 1: ATCA for Telecom

## II. Introduction

This White Paper addresses the problem of network access presented by multiple technologies in today's emerging all-IP network. Led by the explosion of data, in particular IP Video, today's wireless applications demand higher bandwidth from the network. Many approaches are contending to supply user connectivity; Femtocells, UMA, VoLTE, and WiMAX among them. In addition it is clear that TDM and ATM-based solutions have not outlived their usefulness. The problem here is that each of these approaches requires a unique access technology.

Fortunately there are common areas that can be leveraged to provide access via diverse technologies between the Edge and Core Networks from a common ATCA platform. Control plane applications can be hosted by a single application combining signaling interworking as well as integrated Call Control and User Plane management. Many 3<sup>rd</sup> party and in-house solutions address this problem today. What's key for users of ATCA platforms is how to maximize resources such that both development and deployment costs are kept low while allowing for scalability and redundancy.

On the user plane, all access technologies rely on UDP/IP and most are also using Real-Time Transport Protocol (RTP). Interworking to ATM Core Networks is also a requirement in today's hybrid networks. In fact both core and access networks can be either ATM or IP. The details of a femtocell access concentrator solution are a useful case study. Four architectural approaches are outlined, each with its own set of strengths and weaknesses:

1. Single Blade Computers (SBCs) w/ ATCAMezzanineCard (AMC) slots
2. SBC with ATCA Carrier Blades; two options
  - with Local Management Processor (LMP)
  - with Processor AMC card
3. SBC with an ATCA Carrier Blade and Packet Processing Blade
4. Intelligent ATCA Carrier Blades

These solutions are not mutually exclusive. In fact, combined, they represent the range of ATCA's four-part mantra: **Scalability, Flexibility, Redundancy & Cost-Effectiveness**. The goal is to provide the reader with ATCA-based options that will meet his or her particular requirements.

### III. Industry Trends

- IP data is driving telecom network expansion and growth around the world.
- Developed and developing countries are expanding networks and services.
- Economies of scale demand a common platform for new networks and upgrades.
- NEPs must meet market demands for basic, enhanced and advanced services.
- Support is required for new and legacy equipment, technology and services.
- Support flexibility and scalability on different platforms for identical applications without sacrificing reliability or cost-effectiveness.

Today's platforms must be able to adapt and support the mixture of emerging all-IP networks, whether it's 4G, LTE, LTE-Advanced, IMS, NGMN or Legacy technologies. The migration to the all-IP network will not be done in one clean sweep. Even green-field networks must support legacy connectivity. No single technology or vendor can solve all these problems alone. Technological breadth incorporating in-house and 3<sup>rd</sup> party products and partnerships will produce the best-of-breed solutions.

While ATCA has many advantages to NEPs, IBM, Oracle/Sun and HP all offer solutions based on proprietary blade-centric architectures. Perhaps not so ironically each also offers carrier or expansion blades that support standard AMC, PCI or PCIe bus I/O cards. Rack Mount Servers will also continue to play a role in select network applications.

### IV. Diverse Access to the Core Network

Today's networks offer numerous technologies to reach the Core. The key to economical nodes is discovering the commonalities of each access mode and leveraging them into one device. Figure 2 below abstracts these functions to a high level, highlighting functions common to all access technologies. The details of other access technologies and legacy access for 2, 2.5, and 3G systems are available in the Adax Application Notes.

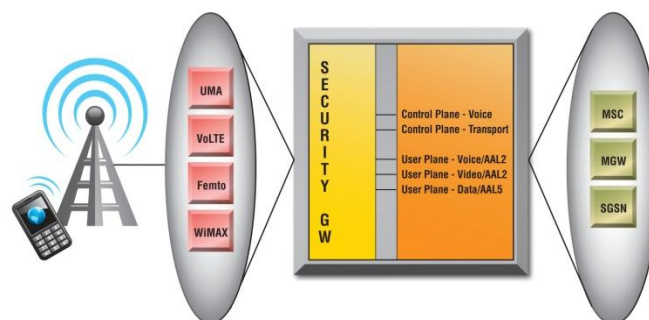


Figure 2: NGMN Generic Access Concentration

## V. FemtoGW/Access Concentration to an ATM Core Network

**A. Functional Requirements:** The FemtoGW/AC for 3G ATM networks requires several functional devices integrated into one system. ATM SS7 Signaling, SIGTRAN Signaling, SS7-SIGTRAN interworking, Transport Signaling for ATM and IP networks, and user plane support for both ATM and IP networks must be present in a FemtoGW/AC. In addition the AC must be secure, requiring the addition of a SecGW and support a redundant, Highly Available software and hardware architecture. Figure 3 shows some of the complexity involved in such a system configuration. AAA and other OA&M functions are left out of this picture in order to focus on the I/O subsystem.

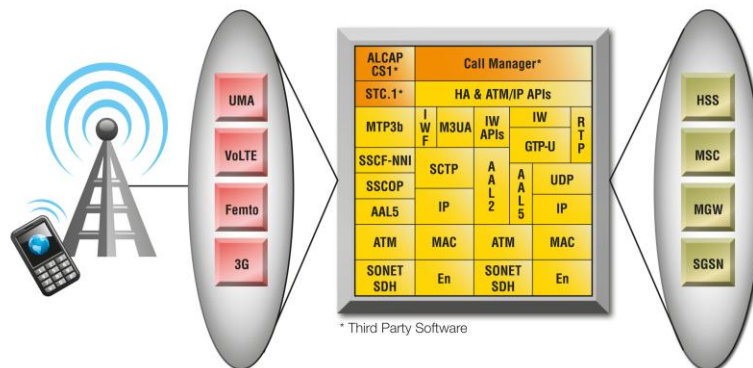


Figure 3: FemtoGW/AC for 3G/ATM Networks

**B. Performance & Capacity Requirements** are set to represent 3 IP-voice Gateways (GWs) feeding traffic into the FemtoGW/AC. It is approximated that at any given time 16,000 users per GW will be active. 3 GWs would yield 48,000 simultaneous users. We assume 1 InterWorking Channel (IWC) per user to support ATM AAL2/CID UDP/IP/RTP IW. At 12.2k bps per voice call we calculate IP packet traffic to reach 608 Mbps for 16,000 users and 1.8 Gbps for 48,000. ATM cell traffic would be 252.67 Mbps for 16,000 and 758 Mbps for 48,000. One ATM-IW card will support one GW for both user and control plane traffic.

### ATM Traffic 766Mbps/ 1 Gbps Unspanned

Packet Type/Size	Header	STF	CID	Payload	Total
Bytes	5	1	3	31	40+13 padding
Bits	40	8	24	244+4 padded	424 (72+248+104)

- Total ATM cell size = 424 bits
- 72b overhead + 244b payload + 4b + 104b padded
- 12,200 bps/244 payload bits = 50 payload packets/call/s
- 50 calls \* 424 bits = 21,200 bps/call
- 21,200 \* 48,000 calls = 1,017,600,000 bps
- Spanning 2 CID payloads across cells:
  - 12,200/324(244+80) payload bits = 37.65
  - 37.65 \* 424 = 15,965 bps/call
  - 15,965 \* 48,000 = 766,340,740 bps

### IP Traffic for 48,000 Users: 1.8Gbps

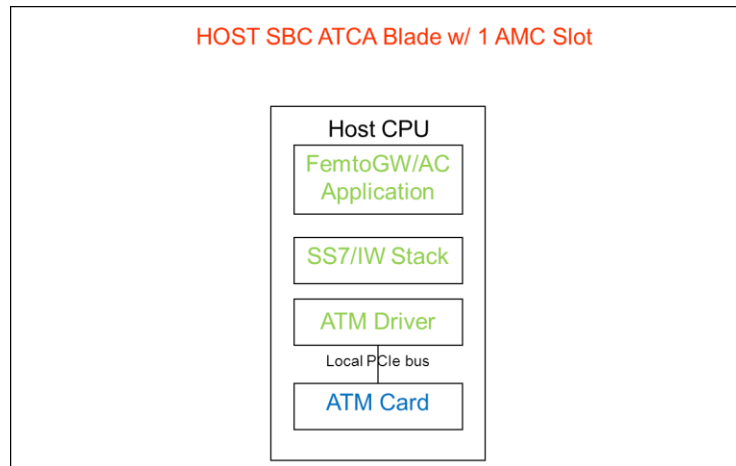
Packet Type/Size	Eth	IP	UDP	RTP	Voice Payload
Bytes	16	20	16	12	31
Bits	128	160	128	96	248

- Total VoIP Packet size = 760 bits
- 512b overhead + 248b payload
- 12,200 bps/244 payload bits = 50 payload packet/s/call
- 50 packets/call \* 760 bits/packet = 38,000 bps/call
- 38,000 \* 48,000 calls = 1.8Gbps (1,824,400,000 bps)

**Figure 4: Detailed Traffic Calculations**

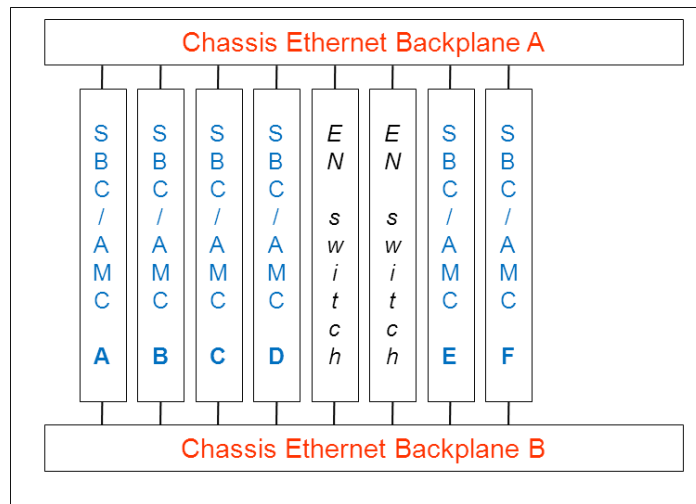
**C. Scalability & Flexibility:** The key to technological and economic success for FemtoGW/ACs is the ability to meet a variety of functional and capacity requirements that can scale up easily and efficiently as capacity needs increase. ATCA is ideally suited to meet this requirement. There are of course, cost/benefit ratios for each option to be considered that will vary widely depending on the network requirements and traffic dimensioning.

- i. **Option 1:** Figure 5 shows the most basic ATCA design possible: One SBC/AMC to host the application and associated control and user plane services. For each initial configuration we are assuming 16,000 simultaneous users before requiring an additional computing resource, in this case the SBC/AMC pair. We then expand that to meet the target requirement of 48,000 simultaneous users in an HA configuration.



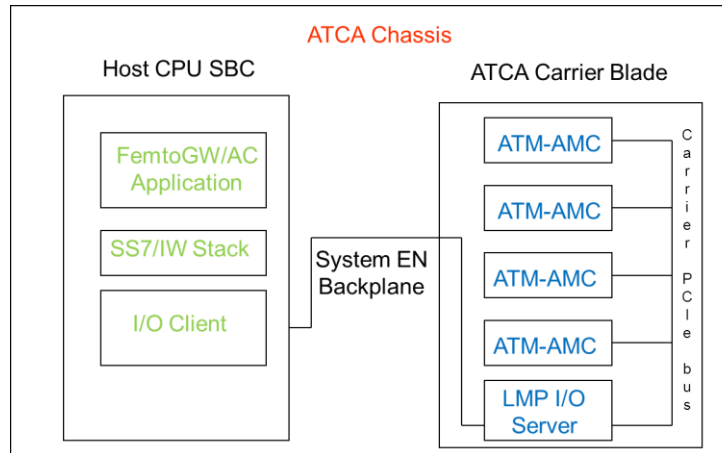
**Figure 5: ATCA Single Board Computer with ATM AMC Card**

In HA this configuration would require two SBC/AMC cards for each 16,000 users. To scale up for each 32,000 users would require 2 pair and 48k 3 pair or 6 SBC/AMCs as depicted in Figure 6.



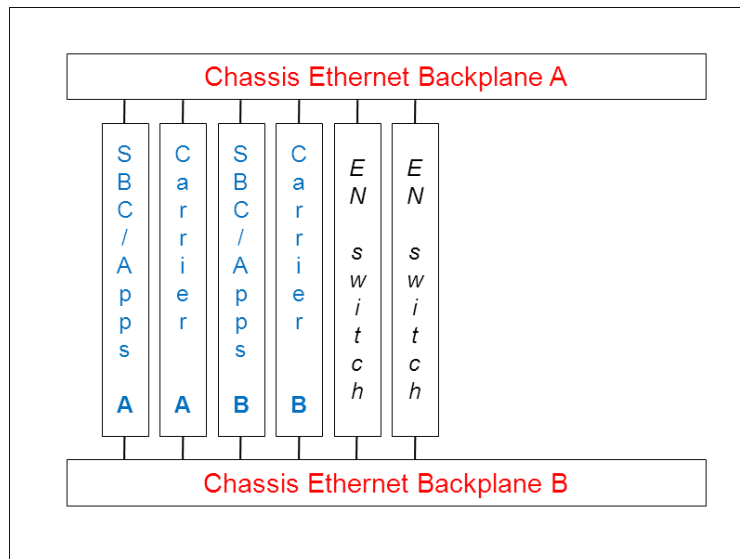
**Figure 6: SBC/AMC HA Configuration for 48,000 Users**

- ii. **Option 2A:** Figure 7 shows a more scalable ATCA design with one SBC dedicated to the FemtoGW/AC application and an ATCA carrier blade to host multiple AMCs. By our calculations this configuration could support three ATM cards each with 16,000 users for a total of 48,000 simultaneous users. One potential limiting factor is RTP processing which must also be done on the host SBC as well as other control plane functions.



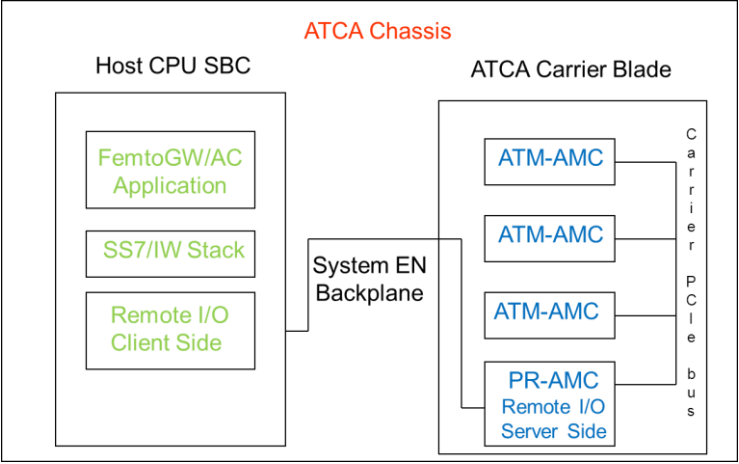
**Figure 7: ATCA Carrier Blade with Local Management CPU**

To support 48,000 users in HA, this configuration would require a total of four (4) blades as shown in Figure 8.



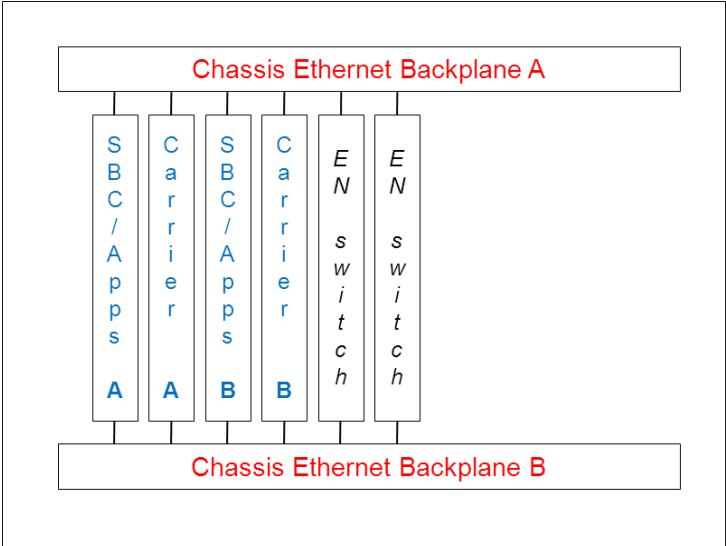
**Figure 8: 48K HA SBC with LMP ATCA Carrier Blade**

**Option 2B:** It is possible to relieve this RTP burden from the Host SBC and push it to the carrier blade by adding a ProcessorAMC as shown in Figure 9, however this requires allocating one AMC bay to the PR-AMC.



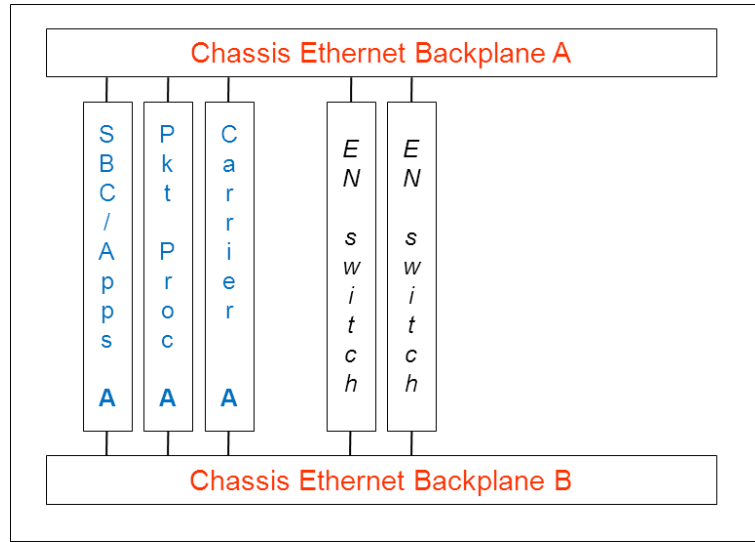
**Figure 9: ATCA Carrier Blade with Local Management CPU**

In HA, this configuration option also requires two pairs of SBC/carrier blades for a total of four slots utilized to support 48,000 users.



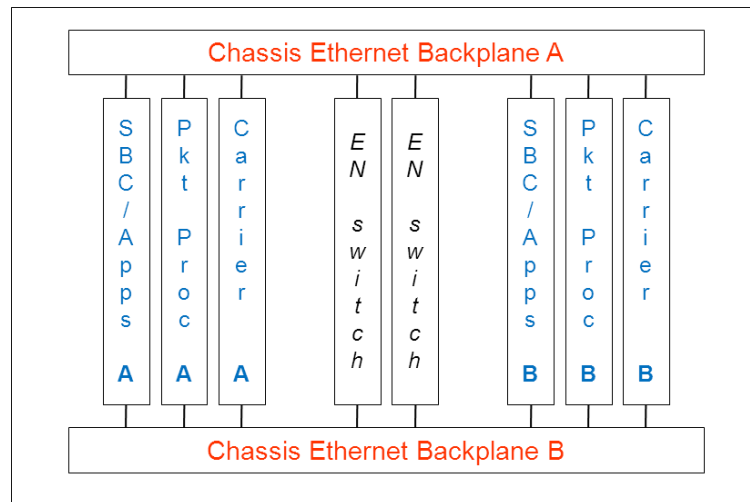
**Figure 10: 48K HA SBC with PR-AMC assisted ATCA Carrier Blade**

- iii. **Option 3:** A configuration that is intended to support more IP-based LTE applications including an ATCA blade dedicated to packet processing is also an option. Control plane applications would reside on the SBC, Fast Path/user plane applications would reside on the packet processing blade and I/O requirements for ATM or GSM I-TDM/Frame Relay Gb would reside on an ATCA carrier blade. Figure 11 depicts this configuration.



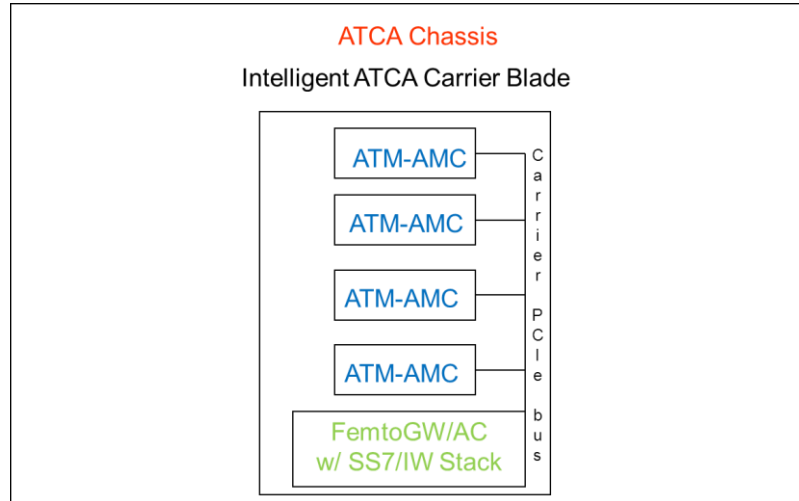
**Figure 11: Option 3 SBC, Packet Processor & ATCA I/O Carrier Blade**

In HA this configuration would require 6 slots; however with the addition of a dedicated packet processing blade, the ATCA carrier would support four ATM-IP AMC cards for a total of 64,000 users rather than 48,000 for the other solutions.

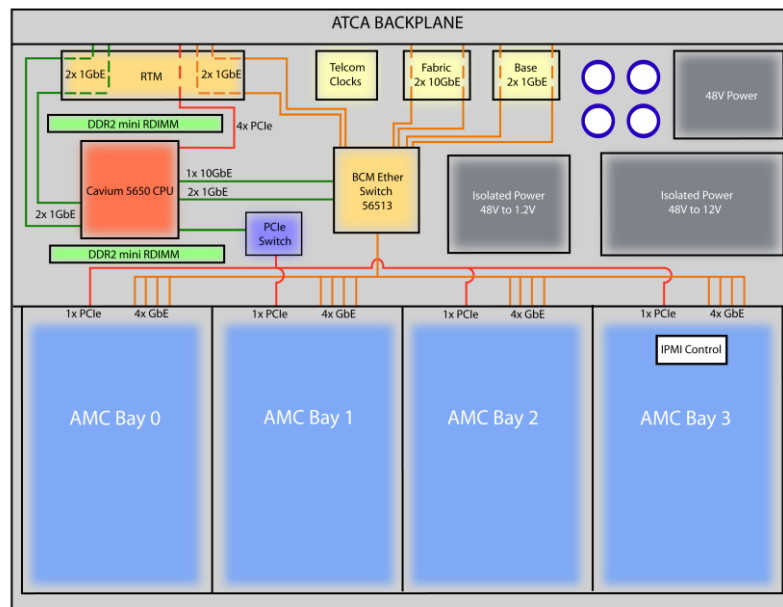


**Figure 12: Option 3 HA: 2x SBC, Packet Processor & ATCA I/O Carrier Blade**

- iv. **Option 4:** Figure 13 shows an example of an intelligent ATCA carrier blade with sufficient NPU/CPU resources to support both control and user plane functions. Figure 14 shows the Adax PacketRunner (APR) as an example of such an intelligent ATCA carrier blade that can host a complete application including RTP or function as an I/O subsystem in the service of a host-based application.



**Figure 13: Intelligent ATCA Carrier Blade**



**Figure 14: Adax PacketRunner Diagram**

## D. Adax Scalability & Flexibility:

- i. **Initial Configuration:** In Figure 15 the FemtoGW/AC on the APR is initially populated with one ATM4 card for ATM signaling and user plane traffic in support of one voice GW. The Adax APR can be populated with up to four AMC cards of like or discrete functions.

The ATM4 card passes signaling traffic to the Cavium-based application which interworks traffic at the MTP3b-M3UA network layer for communication with the IP-based Security GW and femtocell access points.

ATM AAL2/CID user plane traffic is processed on the ATM4 card into UDP/IP packets. These packets are sent to the APR Cavium via dedicated ATM4 Ethernet ports to the Broadcom switch for transport to the Cavium.

Once on the APR Cavium, special features of an Adax application apply the appropriate RTP header which is received from the Call Manager application.

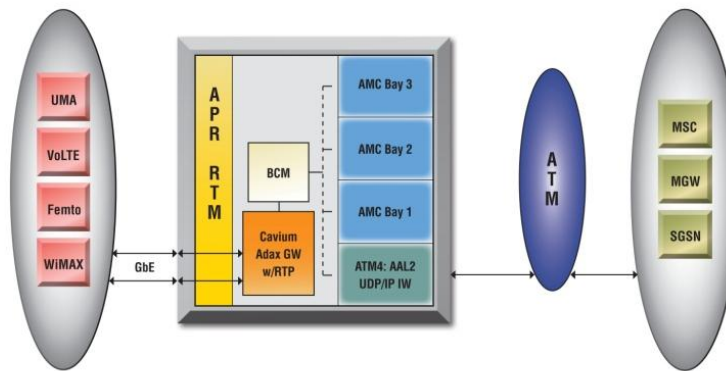
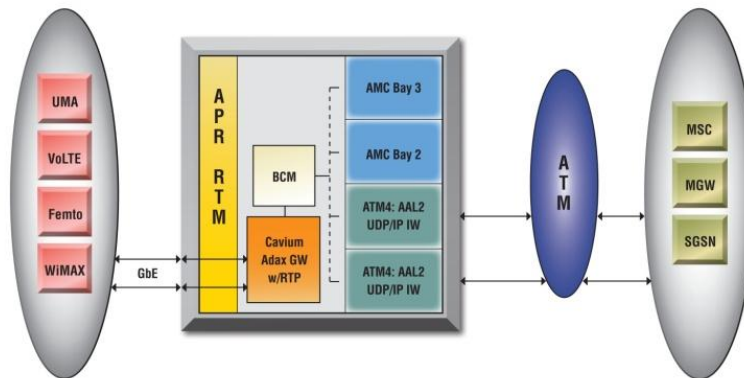


Figure 15: Scalability & Flexibility #1

- ii. **First Level Scaling:** As network needs increase this configuration can be expanded simply by adding one more ATM4 card. Traffic dimensioning figures at the end of this section show that one and very likely traffic from two ATM4 cards can be supported by the RTP application located on the APR Cavium. The economy here is obvious. With careful planning and supported hot swap and hot plug capabilities, application capacity can be easily doubled as shown in Figure 16.



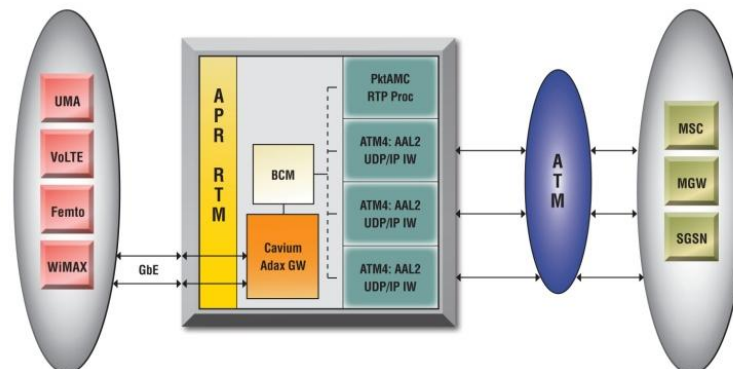
**Figure 16: Scalability & Flexibility #2**

- iii. **Second Level Scaling:** The APR-based node still has room to expand into the two empty AMC bays. Our trials indicate that adding one more ATM4 card and roughly 16,000 more simultaneous users requires additional RTP header processing capabilities. This is accomplished by adding a Cavium-based AMC (PacketAMC) card that will take over RTP functions for all three ATM4 cards shown in Figure 17.

The identical code running on the APR Cavium is simply enabled on the PktAMC. This allows complete flexibility for system expansion as well as re-configuration and allocation of equipment.

User plane traffic must be redirected to the PktAMC from each of the ATM4s. This can be done in a real-time so as to ensure no loss of network traffic during the upgrade.

The APR Cavium is now completely dedicated to the increased signaling and management needs of the application. The addition of the PktAMC also provides the system with four additional 1G Ethernet interfaces to the IP network.



**Figure 17: Scalability & Flexibility #3**

- iv. **Additional Functionality:** Figure 17 shows the addition of the PktAMC into the system. This resource supports other functionality in addition to RTP. Currently implemented stand-alone Security Gateways (SecGWs) can be integrated into each APR configuration or reside on a dedicated APR SecGW blade. This additional capability reduces the overall cost of the system by using fewer computing resources and placing those that remain on a common platform.

**E. Redundancy & High-Availability (HA):** is a requirement for both signaling and user plane traffic. Telecom networks must continue to support five 9's reliability. Many applications such as the Adax SS7-SIGTRAN Signaling Gateway (AdaxGW) support a 1+1 HA architecture. SS7 enhances signaling HA through configuration of links and link sets to the separate elements in a redundant pair. Network Service and SS7 providers have achieved the highest levels of system and therefore network reliability with this architecture. We must apply the same rigor for data integrity to user plane applications. ATCA and associated software applications make it possible to retain the established standard of reliability set by SS7 in the emerging all-IP network.

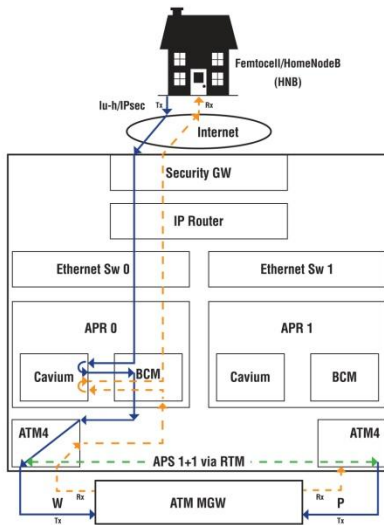
- i. **HA for user plane** redundancy in ATCA is designed in via redundant blades populated with redundant ATM4 cards. The tricky bit is to assure the real-time user plane failover protection, call preservation with minimal data loss and swift recovery times in the event of failures.

Adax believes this challenging task can be accomplished in ATCA by using a combination of Automatic Protection Switching (APS) for optical links and the HA principle of no single point of failure applied to user plane voice calls. With this approach in place one can preserve voice calls, surviving element failures using a failover mechanism between compute elements. These elements include ATM4 card failure on APRs, OC3 APS for single port failure, single ATM4 card failure, APR blade failure or IP Ethernet switch failure. These APS, HA, and failure case recovery features are further detailed below.

- ii. **OC-3 APS:** the user plane application with physical port redundancy. This feature provides for <20ms switching of traffic from one OC3 port to another in 1+1 mode. This redundancy is accomplished via a direct communication link between the *phy* parts on each board. If we take the case of board to board protection and enhance that protection by placing each board on a separate blade we cover three failure modes; line, card, and blade.

As depicted in Figure 18, both the working and protection ports are sending traffic to the Media Gateway (MGW) in the ATM network. The MGW is likewise configured with a 1+1 working and protection port scheme. The MGW sends traffic back to the ATM4 on both ports but the protection port does not forward its data to its host driver. Only the working line establishes bi-directional traffic to and from the ATM network and then on to the Femtocell Home NodeB. Unfortunately this is not enough to preserve the voice call. Both blades and thus ATM4s must know about each other's work

load and be prepared to take over instantaneously when traffic begins to flow on the protection port. Adax HA framework resolves this problem.



**Figure 18: User Plane Protection w/ APS Only**

- iii. **Adax User Plane HA** brings a layer of intelligence to the problem of call preservation. In the event of a line failure on ATM4/APR0, the ATM4 on APR1 must be aware of the ATM-IP interworking channels as well as the current RTP sessions that are active on APR0. The reverse case is true for APR0 assuming a second port on each card is configured for working on APR1 and protection on APR0. Adax User Plane HA ensures that this information is exchanged between APRs and constantly kept up to date. Now each APR is aware of current calls and interworking channels and can take over in case of failure. In addition to the swift failover to ensure call preservation, Adax User Plane HA provides management information to the host application, be that the AdaxGW or 3<sup>rd</sup> party software as depicted in Figure 19.

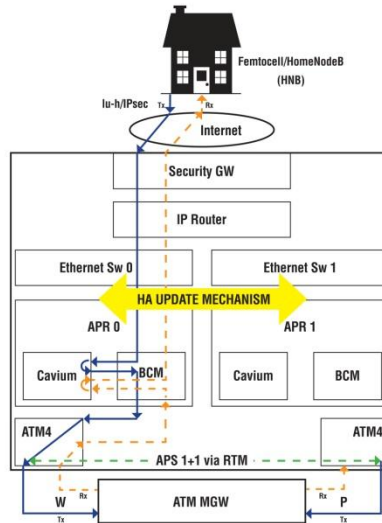


Figure 19: User Plane Protection w/ APS & Adax User Plane HA

- iv. **Failover Scenarios** cover the three cases outlined above. APS is designed first and foremost to protect against line failure. In the simplest case this can be done on a single ATM4 card by configuring port0 as working and port1 as protection. If these ports are connected to separate paths in the optical network then a failure on the working path is detected immediately and the protection port takes over. Since all relevant parameters are either identical or known on the board driver there is nearly no data loss.

In the cases where the OC3 protection port is on another card, or the ATM4 card or host APR blade goes down, Adax User Plane HA must be in play to protect the voice call. Figure 20 shows the failover mechanism at work in any of these cases.

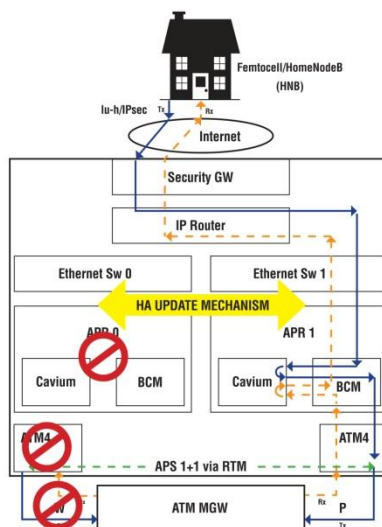
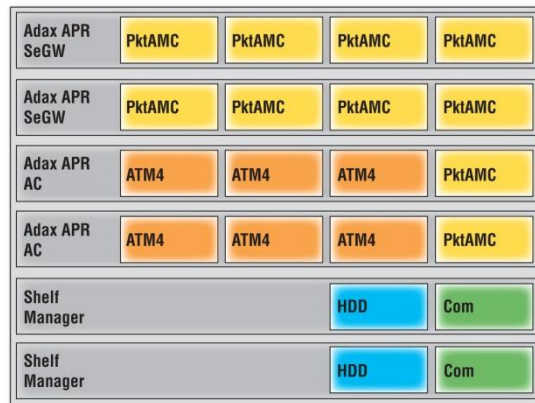


Figure 20: User Plane Protection Failover w/ APS & Adax User Plane HA

**F. Cost-Effectiveness:** ATCA lives up to its promise of scalability, flexibility and redundancy. It also shows that it is a cost-effective architecture for telecom applications. Figure 21 shows the configuration for a fully functional Femtocell Access Concentrator with a SecGW built in. This node could support up to 100,000 simultaneous calls. Since most calls are effectively half-duplex this configuration would offer full redundancy for both user and control plane applications.

Scaling with redundancy is accomplished as depicted below in a 5U, 6-slot chassis. In a comparable cPCI-based solution nearly twice the rack space would be required to support two 4U chassis containing equivalent compute and I/O power as accomplished in this one box solution below.

Should this site require further capacity an additional 5U chassis can be brought in or a new 12U chassis that is 'right-sized' can be installed. The original 5U would then be available for redeployment. In a real-world customer application comparison, this ATCA configuration saves 20-30% over current solutions.



**Figure 21: 5U-6 slot Cost-Effective, Redundant Populated Chassis**

**G. The Adax Advantage** lies first and foremost in the Adax PacketRunner. It is the only ATCA carrier blade on the market with enough computing power to get the job done. Adax AMC protocol controllers running on the APR makes for an unbeatable ATCA combination. Benchmarks comparing the AdaxGW on the APR Cavium with Adax ATM4 and HDC3 AMCs to our current Intel-based cPCI/PMC solutions show a 2.5:1 performance advantage for the identical AdaxGW code running on the APR. There is plenty of horse-power to run the complete signaling and call manager applications while the AMC cards take care of user plane traffic. Adax and ATCA combine to deliver proven, reliable, and cost-effective solutions for your applications.

## VI. Summary & Conclusion

ATCA was designed to deliver flexible, scalable, redundant and cost-effective solutions to the telecom market. Developers and designers have been working with this architecture for several years and are now bringing to market solutions that deliver on the promise.

This example of a FemtoGW/AC clearly illustrates how this key node in emerging all-IP wireless networks can be configured using ATCA. The system is scalable from 2-12U, flexibly transfers functions from one compute element to another, maintains redundancy at all times, yielding a cost-effective solution for network equipment providers and their customers, the network service providers.

In future, the benefits of ATCA will continue to be revealed. All-IP solutions are now being developed and fielded, connecting legacy TDM and ATM network nodes into the IP core. Conversely, IP access devices such as femtocells are being connected to legacy ATM core networks. All of the various permutations can be easily supported in ATCA. Legacy applications can exist natively on ATCA SBCs using virtualized environments and legacy Base Stations can be backhauled using ATCA concentrators. The possibilities are endless, all enabled by ATCA.

Adax offers application developers and system integrators the market's only intelligent ATCA carrier blade and AMC cards capable of hosting the most demanding, complex and sophisticated user and control plane applications for today's complicated requirements. Adax is proud to contribute to the next generation of telecommunications innovation with its current and future ATCA hardware and software products.

## About Adax

Adax, Inc. has a 25+ year history of providing data communications protocols to Tier1 NEPs. We now bring our expertise to bear on the challenges of performance, reliability, and data integrity in the emerging all-IP network. Adax **Foundations for the All-IP Network** includes a full set of hardware and software components for ATCA as well as Rack Mount Servers.

### The Adax ATCA I/O Subsystem

The Adax PacketRunner (APR) is the industry's only intelligent ATCA carrier blade with the power to run applications, not just simply manage the blade. Cavium's Octeon 5650 multi-core processor offers support for the Adax Suite of AMC cards, drivers, and software plus the capacity to host 3<sup>rd</sup> party applications. The Adax PacketRunner combined with the Adax PktAMC for packet processing, ATM4-OC3 for ATM signaling and interworking, and HDC3 for TDM SS7 and ATM provides a complete I/O subsystem for user and control plane applications.



Foundations for the ALL-IP Network

**adax, inc.**  
614 Bancroft Way  
Berkeley, CA 94710, USA  
Tel: +510 548 7047  
Email: [sales@adax.com](mailto:sales@adax.com)  
Web: [www.adax.com](http://www.adax.com)

**adax europe ltd**  
Reada Court, Vachel Road  
Reading RG1 1NY, UK  
Tel: +44 (0)118 9522800  
Email: [sales@adax.co.uk](mailto:sales@adax.co.uk)  
Web: [www.adax.com](http://www.adax.com)

## Appendix A: Glossary

AAL2/CID	ATM Adaption Layer 2 / Channel Identifier
AC	Access Concentrator
AMC	Advanced Mezzanine Card
AMPS	Advanced Mobile Phone System
APS	Automatic Protection Switching
ATCA	Advanced Telecommunications Computing Architecture
ATM	Asynchronous Transfer Mode
cPCI	Compact Peripheral Component Interface
EPC	Evolved Packet Core
HA	High Availability
IMS	Internet Protocol Multimedia Subsystem
IP	Internet Protocol
IWC	InterWorking Control
LTE	Long Term Evolution
MC	Media Converter
MGW	Media Gateway
M3UA	Message Transfer Part Level 3 User Adaptation Layer
MTP3b	Message Transfer Part Level 3 Broadband
NEP	Network Equipment Provider
NGN	Next Generation Networking
PCI	Peripheral Component Interface
PCIe	Peripheral Component Interface Express
RTP	Real-Time Transport Protocol
SAE	System Architecture Evolution
SecGW	Security Gateway
SIGTRAN	Signaling Transport
SS7	Signaling System No. 7
TDM	Time-Division Multiplexing
UDP	User Datagram Protocol
UMA	Unlicensed Mobile Access
VoLTE	Voice Over Long Term Evolution
WiMAX	Worldwide Interoperability for Microwave Access

*The names of companies and products are trademarks or registered trademarks of their respective owners.*