


AppEx Accelerator for Servers

Maximizing Network Performance,
Responsiveness and Availability

AppEx Networks Corporation



Servers

face special technical challenges and networking is a forever topic. Bombarded by the ever-increasing request rate and traffic load, the IT departments struggle to keep up with the demands, adopting the top-notch hardware systems. However, the network performance, responsiveness and availability often lag behind. The gain versus the investment is usually very small.

The reason lies deep in the network protocol itself. With AppEx Accelerator installed in the server systems, it boosts the performance from the protocol level and maximizes the bandwidth utilization. At the same time, it shortens the response delay and makes the connections more robust. All these improvements will provide a lot better user experiences and save the cost in bandwidth and network equipment purchase. As an installable software driver, it offers the lowest cost and shortest time span in testing/trial and deployment, dramatically reducing the risks in infrastructure changes.

Zeta-TCP

Zeta-TCP is the heart-and-soul of AppEx Accelerator. Technical details and analysis can be found in the Zeta-TCP whitepaper. To recap, Zeta-TCP is a series of unilaterally deployable TCP acceleration algorithms. It boosts the TCP performance primarily from the following aspects:

- **Intelligent Congestion Window and Recovery Algorithms**

Zeta-TCP is unique in that it evaluates the network congestions from the variations of both the end-to-end latency and the packet loss rate, as well as some other statistical metrics. Therefore it is capable of more accurately detecting the congestions and their levels. With such prompt and accurate information feedback, it is able to use different Congestion Window and Recovery algorithms to maximize the bandwidth utilization when the conditions are good and fallback to relieve the load of the network when congestions happen.

- **Improved Loss-Detection Accuracy**

Zeta-TCP has an innovative algorithm to very quickly and accurately detect whether a transmitted packet is really lost, especially in the complicated loss scenarios. Hence it is capable of retransmit the lost packets a lot sooner and more accurately compared with the tradition TCP and other variations. Statistics show that with a 5% loss rate, the mistakenly retransmitted packets in Zeta-TCP was only about 1/3 that of the traditional TCP. This means Zeta-TCP wastes a lot less bandwidth than the traditional TCP, in addition to its faster response to the packet losses.

- **Reverse Control and Stimulus**

Zeta-TCP introduces a one-of-a-kind algorithm to stimulate or control the peer's packet sending. This enables it to accelerate the inbound (download) TCP flows. Unlike outbound acceleration, the speed of download is ultimately decided by the peer, i.e., how fast the sender sends. Zeta-

TCP is the only algorithm seen in the industry to encourage and hint the peer to send as fast as it can and accelerate the download.

Engineering-wise, Zeta-TCP is designed to be extremely scalable. It can fit in a variety of systems of different available resources and computational powers. It is also transparent to the protocol stacks so that it can accelerate the traffic without any modifications to the existing protocol stacks. Indeed, the protocol stacks may not even be aware of Zeta-TCP's existence. Unilaterally deployable means Zeta-TCP can be deployed just at one end of the connection and it will be able to accelerate the traffic regardless of what variation of TCP the other end is using. These characteristics are ideal for software distribution.

AppEx Accelerator

AppEx Accelerator is the software implementation of Zeta-TCP and a few other optional components, such as Classifier, Scheduler/Shaper and Compressor, etc. All the optional components can be compiled in or out, tailored to the needs of the target systems.

Portability is one of the important design objectives of AppEx Accelerator. We have isolated a small set of platform wrapper APIs so that for different platforms the work is majorly the infrastructural code and implementation of these common APIs. At the time of writing, AppEx Accelerator supports the following platforms:

- Windows, XP and above.
- Linux, Kernel 2.6.18 and above.
- Windows Mobile 5.0 – 6.5 (WinCE 5.01 – 6.0).

The system resource consumption is configurable. The following table shows the typical numbers for a few different platforms.

Platform	Accelerated Flows	Maximum Memory Usage
Windows server	100,000	256MB
Linux server	100,000	256MB
Windows client	10,000	32 MB
Windows Mobile	1,024	6 MB

Table 1: Memory Usages under Typical Configurations

AppEx Networks Corporation started its Accelerator development in 2006. After years of deployments and refinement, the AppEx Accelerator has reached its maturity. Its outstanding stability, high performance, and unparalleled versatility have made it the best unilateral network acceleration product in the industry.

Gomez Benchmarks

Gomez provides the platforms for Web performance and quality tests which are well recognized by the industry. The following are the test reports of the unilateral installation with AppEx Accelerator

deployed in the web server side. The test was performed over almost 2000 different time slots, accessing from a variety of client sites.

Yellow Line: With AppEx Accelerator
 Blue Line: Without AppEx Accelerator

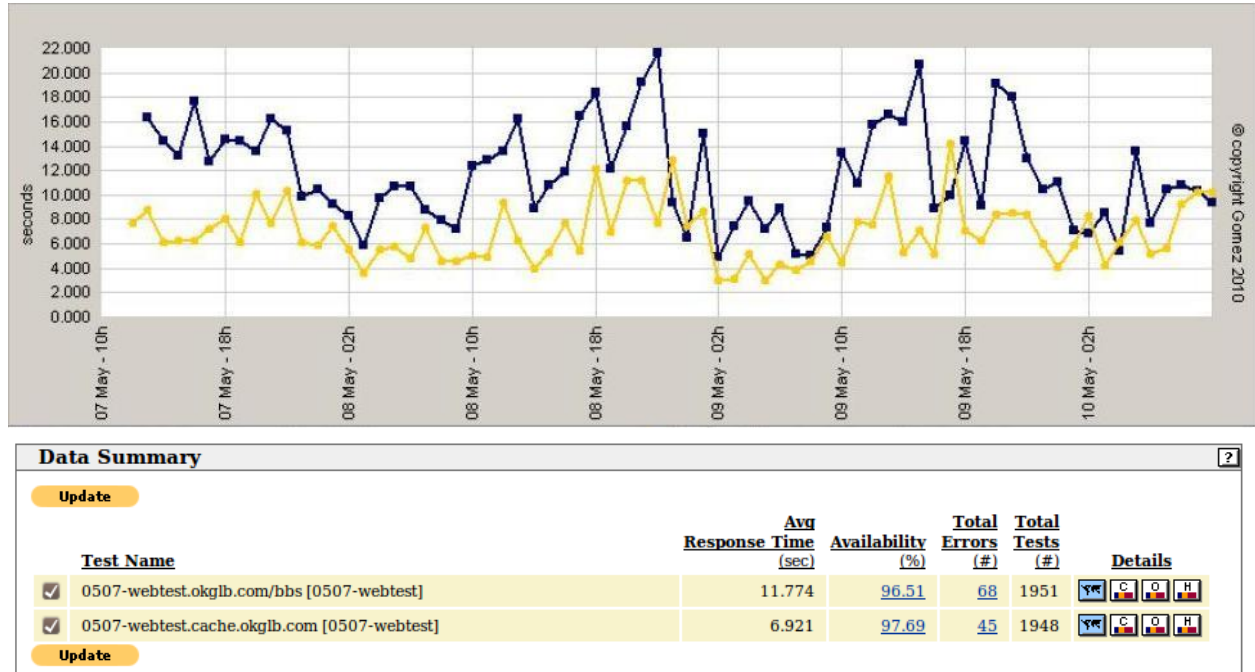


Figure 1: Gomez Test - Dynamic Web Page Loading Time

Gomez Tests	Total Tests	Total Errors	Availability	Avg Response Time
With Acceleration	1948	45	97.69%	6.921 (seconds)
Without Acceleration	1951	68	96.51%	11.774 (seconds)

Table 2: Gomez Test – Statistics

From the test result above, AppEx Accelerator was able to eliminate 34% of web page access failures and shorten the total loading time by 40%.

Where It Can Help

The performance gain of AppEx Accelerator reaches its best in long latency and lossy networks. Such network conditions are often seen in the following scenarios:

- **Cross-operator Communications**

The network trunks between different operators carry a humongous amount of traffic and very often they are overloaded because of the service agreements and the steep charges across the borders of the operators. Hence the traffic usually suffers significant loss and delay.

The cross-operator deterioration is very common in China. From the test results in a few field deployments, for example, the network traffic running between China Telecom and China Netcom usually incurs 4 ~ 5% loss rate and 160+ milliseconds' delay.

- **Lossy Media / Channels**

Wireless / radio based communication channels are inherently lossy due to the characteristics of the Fading Channels. In such media losses happen randomly even without heavy traffic. Wireless 3G is the largest network belonging to this category.

- **Long-distance / International Flows**

Long-distance flows usually have to go through many routers or other network relay devices. Since every one of these routing / relaying devices has limited resources, the more devices the flow has to travel, the higher chances it will incur packet losses. Also longer distance directly or indirectly contributes to the longer delay.

International data traffic, especially Cross-Pacific or Atlantic, also suffers from the distance. On top of that, it almost certainly has to go through different operators.

- **VPN and Other Computation-Intensive Tunnels**

VPN is a mixed scenario. For security it often involves intensive encryption/decryption operations that could take a dramatic amount of computation power. Without expensive specialized hardware device it usually performs a lot slower than normal traffic. The VPN traffic often swamps the gateways with limited resources, causing long delays and high possibilities of losing packets.

To make it worse, the VPN users are very often road warriors, accessing corporate networks from a remote location, across multiple operators, via wireless links. Putting all these together, the VPN connection qualities are often intolerantly bad.

Applications and Real-World Results

Network / Cloud Storage

For this type of applications large amount of data need to be transferred back and forth through the clouds frequently. The network performance is critical to the successful roll out of the storage services. The most majority of the storage services are TCP based, which can greatly benefit from Zeta-TCP's acceleration. Deployed in the storage servers, the Accelerator will not only boost the performance of servicing the users, but also expedite other Cloud tasks such as mirroring and backup.

Corsair Cloud Storage is the mega-storage system developed by the Grid Computing Research Center, Computer Science Department of Tsinghua University. It offers every college student subscriber 2GB private storage and 100GB shared space in the cloud. Corsair Cloud Storage is deployed in CERNET, the

China Education and Research Network (<http://en.wikipedia.org/wiki/Cernet>). The users within CERNET usually have decent speed to access the cloud resources. However, before the deployment of AppEx Accelerator, accesses from the other network operators, such as China Netcom or China Telecom, suffer intolerably slow and flaky data transfers. Since most college employees and off-campus students have Internet subscriptions at home only from operators other than CERNET, the utilization of the cloud is very much limited.

The situation has been dramatically improved after the deployment of the AppEx Accelerator. Many of the cross-operator users' access speed was boosted by an order of magnitude. A few users from different locations with different ISP subscriptions participated in the speed test. The following is part of the test results:

Resource Download (Time: 7/27/2010, 9:30 – 11:00 PM)

ISP	No Accelerator (Kbps)	Accelerated (Kbps)	ISP Bandwidth (Kbps)	Acceleration Ratio
Beijing Netcom	3 (FAILED)	600	512	200
Beijing Netcom	10 (FAILED)	800	2048	80
Beijing Netcom	39	878	1024	22
Beijing Telecom	600	1100	2048	1.9
Feihua Electricity	6.4 (FAILED)	525	1024	80
Northern Telecom	1500	1900	2048	1.2
Great Wall Broadband	105	1300	2048	13

Resource Download (Time: 7/28/2010, 8:30 – 9:30 AM)

ISP	No Accelerator (Kbps)	Accelerated (Kbps)	ISP Bandwidth (Kbps)	Acceleration Ratio
Beijing Netcom	79	628	512	7.9
Beijing Netcom	41	1700	2048	41.5
Beijing Netcom	153	293	1024	1.9
Beijing Telecom	1100	1200	2048	1.1
Feihua Electricity	415	603	1024	1.5
Northern Telecom	1800	1900	2048	1.1
Great Wall Broadband	405	1500	2048	3.7
Huarui Telecom	58	1800	10240	31.0

Table 3: Corsair Cloud Storage Test Result

Online Gaming Services

Robust and fast network connections are crucial for online games. Even the non-realtime MMORPG is now demanding more speed and stability from the networks than ever before. On one hand, the gaming service providers want to attract as many players as possible; on the other more players means more cross-operator and long distance traffic. The gamers' experience forces the gaming providers to invest heavily in improving the network performance.

A common solution to this problem is for the gaming service provider to subscribe to multiple ISPs and host games on all of them. In addition to the risk of steep upfront cost subscribing a broadband dedicated port from a new operator, synchronizing the game states and load-balancing the players also

become big technical obstacles. The load balancer could move a user to a server in a different operator under certain circumstances and deteriorate the gamer's experience a lot. Or the games that has more players from a certain operators than the other may put huge pressures over the game servers in that operator, while the other servers are mostly idle.

AppEx Accelerator offers a great low-budget option for the gaming service providers. With all the servers having the Accelerator installed, the game providers no longer need to be worried so much about players from a different operator or far away locations. Besides all the speed improvement the gamers would enjoy, Zeta-TCP also makes the connections a lot more robust. So the gamers will experience a lot less line-drop rate and more stable ping responses.

Cloud/SaaS, Dynamic Content Servers, Video Servers

Web Content Caching plays an important role in accelerating web browsing and reducing trunk traffic. The caching technology works very effectively for static web contents. However, the recent new trend in Social Media Networks, Cloud / SaaS, etc., swamps the Internet with dynamic web contents. Although many of the child objects of the dynamic content may be static, the delivery of the dynamic object itself is often the bottleneck to load the entire page. Besides, certain applications may not allow the contents to be cached at all for security, privacy or other considerations. The effectiveness of Web Caching is greatly impacted for these scenarios.

Video contents are often not cached for a few reasons. Many web caching services enforce a limit on the size of the objects it caches. The video files are usually too large to be cached. More often, the video content providers streams video into the client-side players. Such video streams are mostly not cached.

With the AppEx Accelerator installed at the dynamic content or other uncacheable content servers, the serving speed of such contents can be dramatically improved. Furthermore, the provider-side deployment is completely transparent to the applications and the clients. The deployment cost is therefore very low. But the user experience improvements would be tremendous.

Corporate Network Resources

The access to the corporate network resources from remote branches or via VPN can greatly benefit from the performance boost of AppEx Accelerator. The connections between the branches or VPN mostly exhibit the characteristics of the long-latency lossy networks, which the AppEx Accelerator is able to make the most of.

For static resources, CDN or Caching solutions can also improve the performance. However, applications such as ERP/CRM and other database accesses generate dynamic data that are not cacheable. Caching can help little in such applications. With AppEx Accelerator software distribution for servers, many of the slow access problems can be easily solved. This can also save the enterprises the cost of deploying the expensive Caching equipments.

AppEx Accelerator is now servicing the ERP applications of Yingjia Group (<http://www.yingjia.cn/>). The company has a 100Mbps subscription from China Telecom. One of its branches is serviced by China Netcom. The branch IPsec tunnels to the headquarter's corpnet. The Accelerator running at the ERP

server side doubled the access speed from the branch, saves the cost of the IT department's original plan of subscribing yet another 20Mbps connection dedicated to the ERP systems.

Video Surveillance Server

Video surveillance devices and servers often employ wireless networks to stream out the video captures. The surveillance video transmissions usually care more about the reliability than the realtimeness, because it is more important to keep as much accurate information as possible. Therefore TCP as transport very well fits the bill for such applications.

Better video quality requires higher bandwidth, which is what the AppEx Accelerator is made for, especially in the flaky wireless networks. What's more, Zeta-TCP's accurate loss detection enables its very fast retransmissions of the lost packets. This makes the video appear a lot smoother.

Some preliminary tests have proved the value of the acceleration. In a 3G video surveillance system, AppEx Accelerate was able to boost the average bandwidth from 494 Kbps to 642.6 Kbps. And the video quality improvement was very apparent.

Conclusions

AppEx Accelerator, designed to be portable, scalable and transparent, is ideal for optimizing the network performance for the server applications. Software distribution is one of the best low-cost channels to offer the server customers the highest value.

Running Zeta-TCP engine under the hood, AppEx Accelerator boosts the network bandwidth utilization, enhances the robustness of the connections, shortens the response time and improves the availability of the server applications. It performs all the optimizations without having to change any applications on the server or the client side. AppEx Accelerator is most effective in the lossy and long-latency network environments.