



## Killer™ and LLR™ White Paper

by

**Harlan 'Tytus' Beverly**

### Executive Overview

On March 22<sup>nd</sup>, 2006, Bigfoot Networks announced Lag and Latency Reduction (LLR™) Technology to the delight of online gaming fans everywhere: but provided no details as to how the technology worked. On July 13<sup>th</sup>, 2006, Bigfoot Networks announced the Killer™ Network Interface Card which is powered by LLR Technology. The purpose of this paper is to discuss how LLR Technology and the Killer NIC work to improve the quality of online game play. This paper is intended for a technical audience, and it is hoped that the reader will come away with an understanding in general of how the technology works. This paper is not intended to be detailed enough to allow copy-cats to reproduce Bigfoot Networks' technology, but rather to allow gamers and enthusiasts to grasp how the technology works in general. All rights and copyrights are reserved, and the technology discussed here is patent pending.

### Competitive Comparison:

There are no consumer products available today that are comparable to Killer. In addition to being a Network Interface Card (NIC), Killer is also a computer within a computer. With Killer's Flexible Network Architecture (FNA™) you can literally run a Linux command prompt in a window on your gaming PC. This command prompt interface is a Flexible Network Application (FNapp™) that will come on the Killer install disc. Killer is the only network card that runs an open-source version of embedded Linux and allows users to write and

download their own FNapps to the card. Killer is also the only consumer level network card that boasts full gaming-network offloading functionality. With Killer, your gaming operating system does not perform any IP or UDP functions or calculations (everything is handled in the card).

Many network products today claim to 'offload' network calculations (like checksum, tcp segmentation, etc.). Those technologies are usually only for TCP/IP networking (which most games that Hardcore Gamers play don't use). Those technologies are also incomplete as they still go through multiple layers of the gaming network stack to eventually get data to the game. With Killer, we completely bypass your gaming PC's operating system and go directly from our card to the game. Our card automatically handles things like IP Reassembly, UDP/IP checksum, UDP and IP header verification and stripping, etc, etc, etc. By bypassing your gaming PC's operating system and allowing Killer to handle everything, Killer can achieve levels of gaming network performance well beyond the offloading features claimed by other consumer networking products (NICs or onboard chipsets).

There are a few network chips/chipsets in the market today that claim the ability to prioritize network packets. Killer also prioritizes network packets, but Killer is the only card that can prioritize both incoming and outgoing packets. However, prioritizing network packets is the least of Killer's features. For these other consumer networking products, prioritization is usually the main or only feature. From a scientific standpoint: prioritizing network packets is a feature best left to routers, because gamers rarely run multiple applications while they are playing games. But with Killer, a gamer can run FNapps while playing games without impacting his or her computer's CPU performance. An example of an FNapp is a file-sharing utility that runs in Killer's Network Processing Unit. This sort of file sharing FNapp would have zero impact to your gaming CPU or your Network: the ultimate in performance!

### Hardware Explained



**Overview:** The Killer NIC has been designed to both offload network processing from the main CPU, and to offer low latency performance (for all Network Traffic). The Killer has specially designed optimizations for User Data Protocol (UDP) traffic, which is the type of packets used in most first person shooter and massively multiplayer game titles. The Killer NIC™ is not a TCP/IP offload engine (TOE), because most games do not use TCP/IP for main game-play. TCP/IP, however, is accelerated by Killer using its normal low-latency NIC design and performing TCP/IP checksum in its Network Processing Unit (NPU). One main difference between the Killer NIC and other network cards (discrete or on-board) is it's focus on low-latency (as opposed to throughput) using video game network processing offload (UDP).

**Board:** Killer NIC is a multi-layered printed circuit board (PCB) design which includes the LLR Technology. LLR Technology is a chip-set that enables the features including MaxFPS™ and UltimatePing™, which will both be explained later.

**NPU:** The first chip of the LLR Technology chipset found on the Killer NIC is a 400Mhz Network Processor Unit (NPU). This chip and the accompanying 64MB DDR Ram are used to run a highly optimized version of embedded Linux. The NPU removes the burden of processing packets from the CPU, as well as processes the packets extremely fast and efficiently. In addition to offloading Network Traffic, the NPU also runs Linux Applications developed within Bigfoot's Flexible Network Architecture (FNA™). More details on FNA can be found in the features section of this technical white paper.

**LLR™ ASIC:** The second chip of LLR is an ASIC (Application Specific Integrated Circuit) that contains proprietary features that work with the NPU to offer the features of LLR (to be discussed later).

**PCI:** PCI 32-bit was selected for three reasons:

1. PCI 32-bit is capable of tremendous throughput capabilities approaching gigabit speeds in unshared PCI buses. Even in shared PCI buses, PCI 32-bit's throughput capabilities are well in excess of the needs of any modern game, and tens of magnitudes greater than a Cable or DSL link could possibly saturate. Most games in 2006/2007 use much less than 1Megabit per second of throughput.
2. The latencies of PCI and PCI-E are similar, if not slightly better in PCI (unshared configuration). Latency is critical in nearly all games.
3. PCI-E slots are used in many gaming systems for graphics, but not all of those motherboards have any free PCI-E slots for expansion (and instead rely on PCI). In addition, some AGP motherboards have no PCI-E slots at all. Almost all motherboards have a PCI 32-bit bus however, and we want the majority of gamers to have access to the benefits of the Killer™ NIC.

#### **Gigabit Connectivity:**

Gamers expect gigabit connectivity, but they do not usually need full gigabit throughput. Gigabit connectivity is still good for gamers because it allows packets to travel across local network hops faster, meaning better throughput, and faster latency times. Even if this is only in a gamers LAN, every millisecond counts! Killer also supports 10/100 Megabit connectivity for older routers and cable modems.

#### **USB 2.0 Port:**

The USB port included on Killer is available to the Linux core running on Killer's NPU. The port is not available as an extra USB port for the host, which cannot see it or use it except via an FNapp. Devices that attach to the USB port are available to the Linux core and users who write FNapps™ (See FNA below). This allows USB Keys to hold file sharing data in file sharing FNapps... allows USB

# KILLER™

NETWORK.INTERFACE.CARD

microphones or headsets for use in VOIP FNapps... the list is endless. Anything you could connect with a USB Key, someone (or you yourself) can write an FNapp to use it with your network, without impacting your host CPU or even using the PCI bus at all! (see more below).

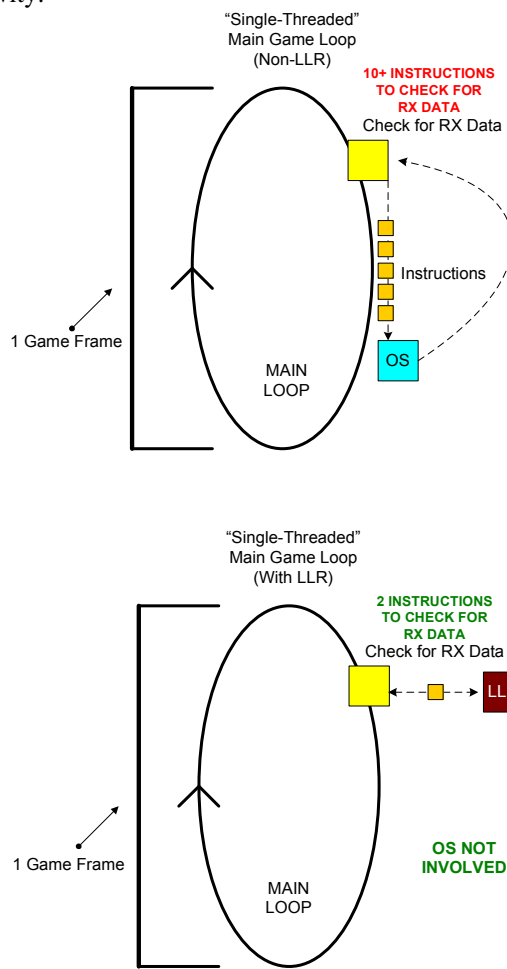
## Technologies Explained

**GameFirst™:** When LLR is enabled, the NPU classifies inbound and outbound packets based on their type. Gaming packets are determined by the UDP packet type, and then are processed in real-time by the NPU and **ONLY** the data within the UDP packets is placed directly to or taken directly from the game's program space. These operations are given a higher priority than other network traffic.

**MaxFPS™:** MaxFPS™ will increase the Frames Per Second (FPS) in most gaming systems. It does this by reducing the CPU utilization due to networking, and speeding up the main game loops of the game (even when no network traffic is present). For gaming systems that have older graphics cards, the additional performance in CPU, cache, and main system memory will improve the efficiency of the older graphics card allowing it to run at more FPS, or at higher resolutions and settings. For gaming systems that have newer graphics cards, the FPS performance is usually limited by the performance of the main gaming loop or the CPU's ability to get data to the card (memory bottlenecks). MaxFPS™ will improve the speed of the main gaming loop and reduce the CPU utilization and main system memory thrashing thus improving FPS, or allowing higher resolutions and settings.

Most games today are mainly single-threaded designs, or are multi-threaded but are controlled by a main game loop. When a main game loop is single threaded (or networking is performed inside the main game loop) (*figure 1*), it will usually poll with non-blocking receives or call non-blocking select on a socket EVERY TIME THROUGH THE MAIN GAME LOOP, regardless of if there is actual network activity or not. This means that many levels of the network stack are traversed to determine if data

is present or not, even when there is NO network activity.



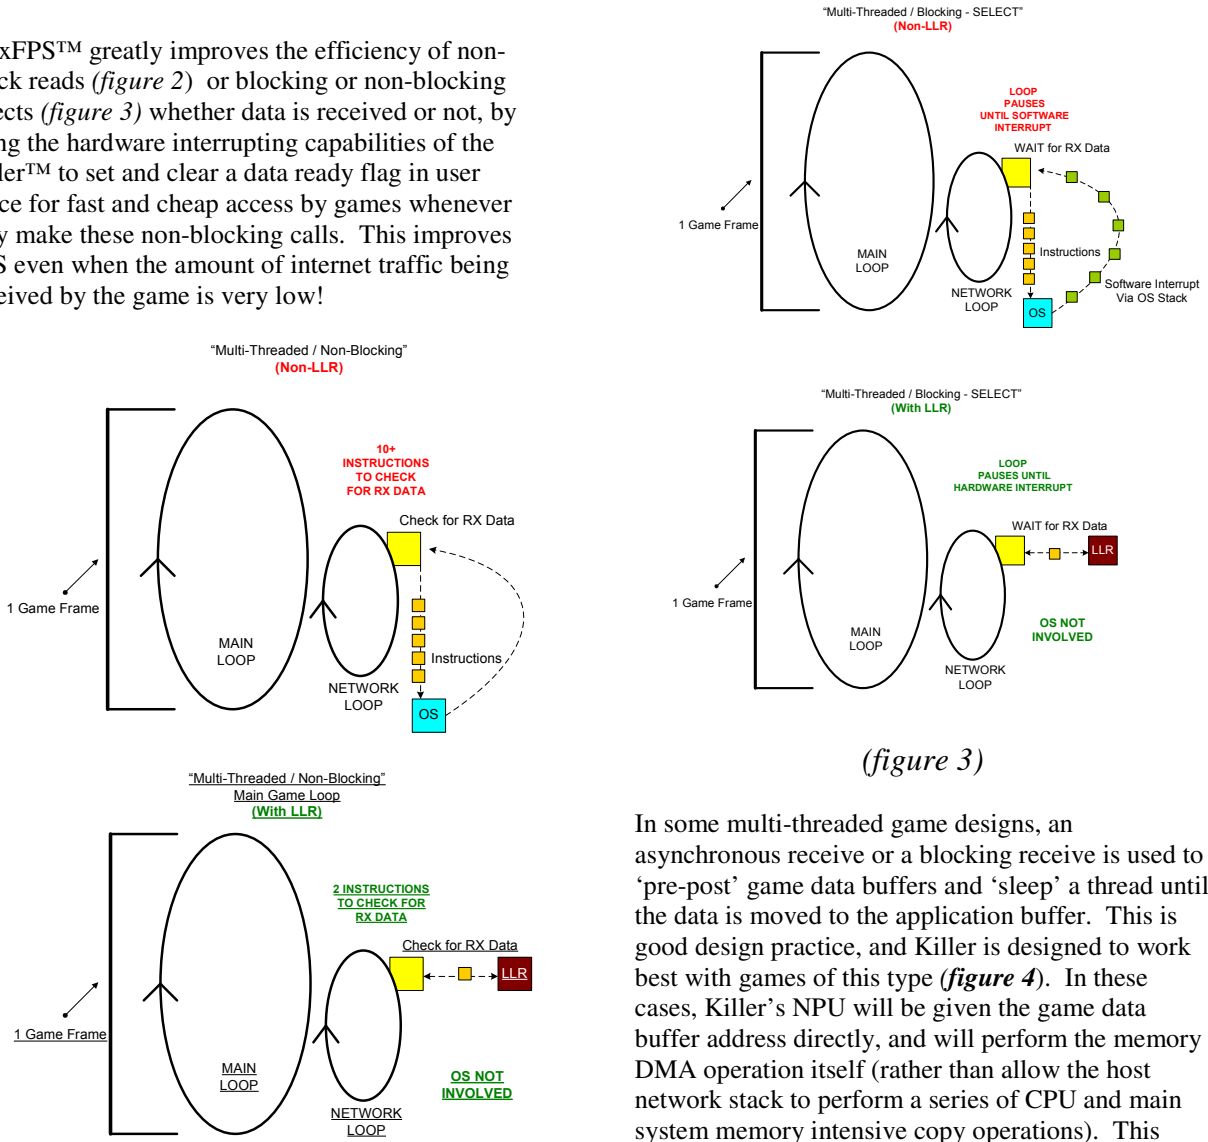
(figure 1)

MaxFPS™ alleviates CPU, cache, and main system memory bottlenecks by decreasing the CPU utilization of single threaded main game loops that poll with non-blocking receives or call the socket command select, and by decreasing the CPU utilization of multi-threaded game loops that poll or call select or perform blocking or asynchronous reads. Whenever CPU utilization is decreased an equal decrease in cache and Main system memory is usually achieved.

# KILLER™

NETWORK.INTERFACE.CARD

MaxFPS™ greatly improves the efficiency of non-block reads (figure 2) or blocking or non-blocking selects (figure 3) whether data is received or not, by using the hardware interrupting capabilities of the Killer™ to set and clear a data ready flag in user space for fast and cheap access by games whenever they make these non-blocking calls. This improves FPS even when the amount of internet traffic being received by the game is very low!



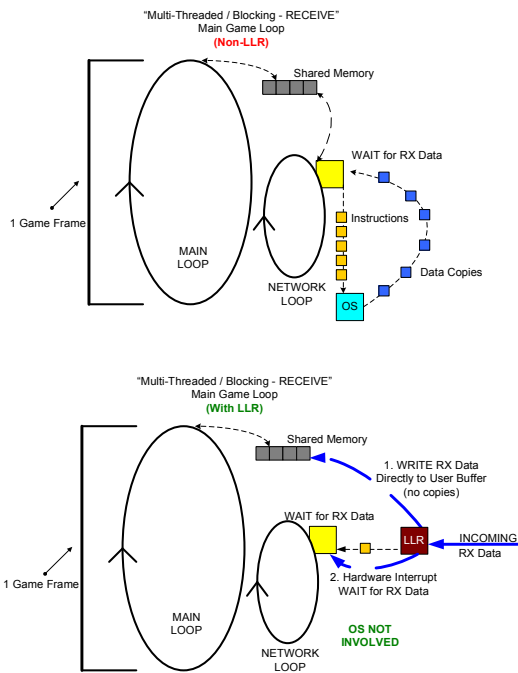
(figure 3)

(figure 2)

In some multi-threaded game designs, an asynchronous receive or a blocking receive is used to 'pre-post' game data buffers and 'sleep' a thread until the data is moved to the application buffer. This is good design practice, and Killer is designed to work best with games of this type (figure 4). In these cases, Killer's NPU will be given the game data buffer address directly, and will perform the memory DMA operation itself (rather than allow the host network stack to perform a series of CPU and main system memory intensive copy operations). This reduces CPU overhead as well as reduces the amount of data that must flow into and out of main system memory.

# KILLER™

NETWORK INTERFACE CARD



(figure 4)

**UltimatePing™:** This will reduce the effective UDP-internal ping time (ping using UDP sockets). A UDP-internal ping is very different than a standard ICMP ping such as one might get when calling 'ping' from a command line. UDP-internal ping runs over a UDP socket and includes such overhead as UDP checksum buffer copy, real data transfer, stack efficiency (including number of copies), etc. Further, a UDP-internal ping (ping/timestamp on top of UDP) uses MaxFPS technology when running through Killer. See Raw Performance Section below for more details.

UltimatePing™ is based on 2 technologies. First, UltimatePing™ ensures that neither your operating system nor your network card are introducing unnecessary latency into your game. UltimatePing™ does this by ensuring that interrupts from the NIC and within the operating system are efficient and NOT delayed, especially for gaming traffic. UltimatePing™ makes sure EVERY packet

(especially a gaming packet) is considered urgent information.

Second, UltimatePing™ uses a side-effect of MaxFPS™ technology to improve the response time in the main game loops or threads. This beneficial side-effect of MaxFPS™ technology is its ability to interrupt the host on a receivable event and set a data ready flag that exists close to the main game loop (user space). Polling or selecting on that ready flag is much faster, and data ready indication arrives at interrupt speeds directly to the host app. Further, if MaxFPS™ uses its direct transfer of data to host app memory, Latency is further reduced by eliminating the time for extra copies.

**PingThrottle™:** PingThrottle™ is a user-controlled setting that literally adds latency to any outbound network traffic, effectively increasing the effective ping a gamer has to a server. This is helpful when you are hosting a server on your LAN that other players are connecting to via the WAN (broadband/etc), and you would like to increase your ping to maintain a standard of fair play. Another use for this is to handicap your system in order to train and hone gaming skills.

Some gamers have been known to 'cheat' by abnormally raising ping to say 400ms and causing a game to reduce accuracy. This cheat is not possible with PingThrottle™ for 2 reasons:

- 1.) The max latency that can be added is only about 20ms (give or take).
- 2.) The adjustment is done in the NPU itself rather than in software, so the added latency is real.

Obviously not all gamers will need this! But it's there, just the same.

**FNA™:** FNA stands for Flexible Network Architecture. It is in effect, the infinite flexibility of Killer. For the average user, this means there is and will be FNapps (utilities/programs/game enhancers/etc.) that can be downloaded from [www.KillerNIC.com](http://www.KillerNIC.com) and run on your KillerNIC.

# KILLER™

NETWORK.INTERFACE.CARD

FNapps™ are designed to allow a user to run an application with a minimal or reduced impact on the main system's CPU, Memory Subsystem, Caching, Hard-Disk, etc. These FNapps can be anything from simple packet monitoring utilities [like fire-walls, etc] to full-blown VOIP programs or file-sharing systems: even mini-game servers/chat servers. FNapps can be designed to make use of the dedicated USB port as well as the gigabit Ethernet port. In addition, FNapps integrate easily (and can communicate easily) with host apps (for gui interfaces/etc.).

Gamers interested in programming, Linux gurus, other users (like web-servers, investment bankers, etc.), Bigfoot Partners, games developers, gamers...anybody who wants...may develop FNapps (applications that run with FNA). Bigfoot Networks will make available the download of a patched open source Linux compiler and Linux source code. This is a standard embedded Linux build, with minor modifications (also open source) to support specific RAM and other hardware configurations of our device. With these tools and the ample amount of documentation and examples (including MS Windows interfacing examples) of how to create FNapps, it is very easy and fast to make one. As part of a host (Windows) program, a programmer may detect that Killer is present, and optionally download the FNapp to the card (or not if it's not present). No special downloader or software is required by FNapps, which can build it into their host program directly. Communicating with Killer from a host is also made very simple for a programmer, and requires no special driver (other than LLR™'s normal drivers). FNapps will give gamers an edge, and enable countless other concepts and technologies to be developed.

## The Author

Harlan Beverly is a Mad Scientist and a CEO: a rare combination. He combined his passion for online gaming, entrepreneurship, and engineering skills in Network Acceleration to help found Bigfoot

Networks, Inc in 2004. Bigfoot Networks has since raised over \$4 Million and won numerous awards for its business plan, including a \$100,000 top prize in the University of Texas MOOT Corp. competition. Harlan honed his business and technical skills while working at Intel, where he filed over 19 patents in the field of Network Acceleration. As an Architect at Intel, Harlan was responsible for the complete solution, including software, hardware, and business cases, for corporate server networking products. After leaving Intel, Harlan joined a start-up called Britestream Networks where he helped develop the first 100% CPU off-load security solution that combines SSL processing with TCP/IP processing on the same chip. Harlan received his MBA from the University of Texas in 2004, and his B.S.E.E from Ohio Northern University in 1999. Lag is a very personal and real problem to Harlan, because he and his wife met while playing Ultima Online in the 90's. When not working or playing with son Christopher or daughter Samantha, you are likely to find Harlan at his ranch tending to his horses, screaming at Lag while playing online games, or tinkering with his multimedia room.

## More Information

Harlan Beverly is the President and CEO of Bigfoot Networks, Inc. Bigfoot Networks is working to develop products that address today's lag problems in innovative new ways. He can be contacted as follows:

Mr. Harlan Titus Beverly, CEO  
Bigfoot Networks, Inc.  
3925 W. Braker Lane, 3rd Floor  
Austin, TX 78759

Web Site: <http://www.BigfootNetworks.com/>