

User Guide

EVGA X299 FTW-K

Specs and Initial Installation

Table of Contents

User Guide	- 1 -
EVGA X299 FTW-K	- 1 -
Specs and Initial Installation.....	- 1 -
Before You Begin... ..	- 4 -
Parts NOT in the Kit.....	- 5 -
Intentions of the Kit.....	- 5 -
Motherboard Specifications.....	- 6 -
Unpacking and Parts Descriptions	- 8 -
Intel X299 FTW-K Motherboard LED reference	- 9 -
Intel X299 FTW-K Motherboard Component Legend	- 12 -
PCI-E Slot Breakdown	- 24 -
M.2 and U.2 Slot Breakdown	- 25 -
Installing the CPU	- 26 -
Installing the CPU Cooling Device	- 27 -
Installing System Memory.....	- 28 -
Installing the I/O Shield and I/O Cover	- 29 -
Installing the Motherboard.....	- 29 -
Securing the Motherboard into a System Case	- 30 -
Installing M.2 devices.....	- 32 -
Tested CPU.....	- 35 -
Tested U.2.....	- 35 -
Tested M.2 Key-M	- 36 -
Tested M.2 Key-E.....	- 36 -
Connecting Cables.....	- 37 -
Onboard Buttons.....	- 47 -
First Boot	- 48 -
M.2 SSD, PCI-E SSD, and NVMe SSD Installation steps	- 50 -

Internal RAID Controller	- 52 -
Fan Header DC and PWM setup	- 87 -
Setting Up SLI and PhysX.....	- 91 -
Realtek HD Audio Manager	- 95 -
Configuration of the Killer NIC Software	- 113 -
Installing Drivers and Software	- 131 -
Windows 10 Driver Installation.....	- 131 -
Warranty and Overclocking.....	- 132 -
Troubleshooting	- 133 -
Replacing a BIOS chip	- 133 -
SSD / HDD is not detected	- 136 -
System does not POST, and POST code indicator reads “C”	- 138 -
System does not POST, and POST code indicator reads “55”	- 139 -
System does not POST, and POST code indicator reads “d7”	- 139 -
Have a question not covered above, or want some online resources?	- 140 -
POST Beep codes	- 141 -
POST Port Debug LED	- 142 -
POST Codes	- 143 -
EVGA Glossary of Terms	- 148 -
Compliance Information	- 151 -

Before You Begin...

EVGA welcomes you to the next generation of Intel Enthusiast performance: the X299 FTW-K! The X299 platform supports the newest Skylake-X and Kaby Lake-X processors. EVGA X299 motherboards further refine high-performance with multiple options for all the latest SSD options with support for U.2, M.2 and PCI-E drives and is Intel® Optane™ Memory Ready – a revolutionary higher-density memory interface, based on 3D XPoint Technology, delivers a new generation of SSDs designed to obliterate loading times for gamers. However, the X299 platform also supports all the features you've come to expect from EVGA, including up to 128GBs of Quad-Channel DDR4 memory at maximum memory speeds of up to 4000MHz+ for Skylake-X, and 4133MHz+ for Kaby Lake-X (OC), Gigabit-NIC support, USB 3.0 and USB 3.1 Type-A and Type-C support, an updated UEFI\BIOS GUI, PWM fan control and a variety of SATA options to fit everyone's needs. The X299 FTW-K is built with a 8-layer PCB, featuring a CPU socket with 150% higher Gold content powered by an Advanced 14-phase Digital VRM (10 Phase VCore, 4 Phase Memory PWM), providing industry-leading stability for all your applications.

Furthermore, this board is designed not ONLY for overclockers, but also for gamers with NVIDIA® 3-Way SLI + PhysX Support without the need for PLX chips, blazing-fast networking featuring the Killer E2500 NIC and an Intel i219, Dual M.2 Key-M, Dual U.2, 8 SATA 3/6g and much more!

Lastly, a motherboard is only as good as its BIOS, and the EVGA X299 FTW-K features an updated UEFI\BIOS GUI with a focus on overclocking and functionality in a lean, straight-forward package. You won't need to be an expert to configure your motherboard, but if you are, you'll find features unavailable anywhere else.

Combining the best of current technology with the latest innovations, EVGA is further refining motherboard performance!

Parts NOT in the Kit

This kit contains all the hardware necessary to install and connect your new EVGA X299 FTW-K Motherboard. However, it does NOT contain the following items, which must be purchased separately in order to make the system fully-functional and install an Operating System:

- ❑ **Intel Socket 2066 Processor**
- ❑ **DDR4 System Memory**
- ❑ **CPU Cooling Device**
- ❑ **PCI Express Graphics Card**
- ❑ **Power Supply**
- ❑ **Hard Drive or SSD**
- ❑ **Keyboard / Mouse**
- ❑ **Monitor**
- ❑ **(Optional) Optical Drive**

EVGA assumes you have purchased all the necessary parts needed to allow for proper system functionality. For a full list of supported CPUs on this motherboard, please visit www.evga.com/support/motherboard

Intentions of the Kit

When replacing a different model motherboard in a PC case, you may need to reinstall your operating system, even though the current HDD/SSD may already have one installed. Keep in mind, however, you may sometimes also need to reinstall your OS after a RMA even if your motherboard remains the same due to issues that occurred prior to replacing the motherboard.

Motherboard Specifications

- ❑ **Size:**
EATX form-factor of 12 inches x 10.375 inches (304.8x276.7mm)
- ❑ **Microprocessor support:**
Intel Socket 2066 Processor
- ❑ **Operating Systems:**
Supports Windows 10 64bit only
- ❑ **System Memory support:**
Supports Quad Channel DDR4 up to 4000MHz+ Skylake-X, and 4133MHz+ Kaby Lake-X (OC).
Supports up to 128GB of DDR4 memory.
- ❑ **USB 2.0 Ports:**
4x from Intel X299 PCH – 4x internal via 2 FP headers
Supports transfer speeds up to 480 Mbps with full backwards compatibility
- ❑ **USB 3.0 Ports:**
8x from Intel X299 PCH – 6x external, 2x internal via 1 FP headers
Supports transfer speeds up to 5Gbps with full backwards compatibility
- ❑ **USB 3.1 Ports:**
2x from ASMedia ASM2142 – 2x external, 1x Type-C, 1x Type-A
Supports transfer speeds up to 10Gbps with full backwards compatibility
*Type-C will require a Type-A adapter for backwards compatibility
- ❑ **SATA Ports:**
Intel X299 PCH Controller
8x SATA 3/6G (600 MB/s) data transfer rate
 - Support for RAID0, RAID1, RAID5, AND RAID10
 - Supports hot plug
- ❑ **Onboard LAN:**
1x Intel i219v Gigabit (10/100/1000) Ethernet
1x Killer E2500 Gigabit (10/100/1000) Ethernet MAC+PHY
Ethernet Teaming NOT Supported
Killer Doubleshot Pro supported with Killer Wireless AC 1535 or 1525
* Killer Wireless AC is **NOT** included and must be purchased separately

❑ **Onboard Audio:**

Realtek Audio (ALC1150)

Supports 8-channel (7.1) audio with Optical S/PDIF Out

❑ **Power Functions:**

Supports ACPI (Advanced Configuration and Power Interface)

Supports S0 (normal), S3 (suspend to RAM), S4 (Suspend to disk - depends on OS), and S5 (soft - off)

❑ **PCI-Express Expansion Slots:**

4x PCI-E x16 slot 2x16, 2x8*

1x PCI-E x4 slot

1x PCI-E x1 slot

*LANES PER SLOT CAN VARY BASED OFF OF CPU, PLEASE SEE PAGE 24 FOR LANE BREAKDOWN BASED OFF OF CPU.

❑ **PCI-E 3.0 Support:**

Low power consumption and power management features

❑ **SLI and Crossfire Support:**

3-Way SLI

3-Way Crossfire

❑ **Additional Expansion Slots:**

1x M.2 Key-M 110mm slot PCI-E/NVMe

1x M.2 Key-M 80mm slot PCI-E/NVMe & SATA/Optane

Key-M 80mm Supports Intel Optane

1x M.2 Key-E slot (via vertical adapter)

2x U.2 slot

❑ **RGB LED Header:**

4x 4-pin RGB LED header (For included RGB lit covers ONLY)

❑ **VROC Header:**

1x 4-pin VROC header

❑ **Fan Headers:**

2x 4-pin PWM controlled headers

5x 4-pin DC\PWM headers

-3pin(DC)\4pin(PWM) modes are selected and static RPM set in BIOS

ALL FAN HEADERS HAVE A MAXIMUM POWER OF 1 AMP @ 12 VOLTS (12 WATTS) EXCEDING THIS LIMIT WILL CAUSE THE BOARD IRREPARABLE DAMAGE.

Unpacking and Parts Descriptions

The following accessories are included with the EVGA X299 FTW-K Motherboard:



The EVGA X299 FTW-K Motherboard

This motherboard contains the Intel X299 chipset.




I/O Shield

Installs in the system case to block radio frequency transmissions, protect internal components from dust, foreign objects, and aids in proper airflow within the chassis.



2x SATA 3G/6G Data Cables

Used to support the SATA protocol and each one connects a single drive to the motherboard.



Thermal Pad for M.2

Thermal Pad for M.2 Key M SSD's.



Vertical M.2 Adapter

Connects the WIFI card to the vertical M.2 adapter.



I/O Cover

This optional cover attaches to the PCB and covers the I/O area



USB 2.0 Bracket (4 port)

Connects to 2 9-Pin USB 2.0 headers to the rear of the chassis.



RGB Audio Cover

RGB LED audio/trace cover



Case Badge

EVGA enthusiast Case Badge*



Installation DVD

Contains drivers and software needed to setup the motherboard.




User Manual

The user manual you are reading right now!

* Decals, labels and case badges MAY vary over time, and the images used are merely reflective of the decal/label/badge designs typically included.

Intel X299 FTW-K Motherboard LED reference

The EVGA X299 FTW-K Motherboard has several LEDs indicating power, connectivity, and activity. Below is the location of the LEDs and their function.



LED Legend		
1. Power Button	7. CATERR	13. CC_LED
2. Reset Button	8. VSM	14. PE2 Status/Mode
3. POSTCODE Indicator	9. VCC	15. PE3 Status/Mode
4. Removable BIOS LED 1	10. 5VSB	16. PE4 Status/Mode
5. Non-Removable BIOS LED 2	11. SKX Detect	17. PE5 Status/Mode
6. FIVR	12. KBX Detect	


1. Power Button
 - a. RED: System is on
 - b. GREEN (blinking): System in Sleep mode
2. Reset Button
 - a. WHITE: blinks with storage access
3. POST Code indicator
 - a. After bootup, this will display the CPU temp.
 - b. During boot it will cycle many different hexadecimal post codes with a range of 00-FF and this indicates what aspect of the Power On Self Test (POST) is currently running.
 - i. *For list of POST Codes, please see Page 142.
4. Removable BIOS LED
 - a. RED: Active BIOS Chip (only 1 will be lit at a time)
 - b. This LED indicates the removable BIOS chip is active
5. Non-Removable BIOS LED
 - a. RED: Active BIOS Chip (only 1 will be lit at a time)
 - b. This LED indicates the non-removable BIOS chip is active
6. FIVR
 - a. RED: Indicates CPU integrated voltage regulation failure.
 - b. OFF: No issues detected on FIVR
7. CATERR
 - a. CATERR stands for Catastrophic Error on the processor.
 - b. RED: Processor error has occurred.
 - c. Off: No error state detected in the CPU.
8. VSM
 - a. YELLOW: Voltage detected (**Does not mean PSU is outputting in-spec, only that this specific voltage is detected**)
9. VCC
 - a. RED: Voltage detected (**Does not mean PSU is outputting in-spec, only that this specific voltage is detected**)
10. 5VSB
 - a. WHITE: Voltage detected (**Does not mean PSU is outputting in-spec, only that this specific voltage is detected**)
11. SKX LED
 - a. YELLOW: Indicates a SkyLake-X processor is installed into the socket.
12. KBX LED
 - a. WHITE: Indicates a Kaby Lake-X processor is installed into the socket.

13. CC_LED
 - a. USB 3.1 Type-C detection.
 - b. RED: Indicates device is attached and detected to USB Type-C.
 - c. OFF: No device detected.
14. PCI-Express LED for PE2. The LED will remain off when this PCI-E slot is disabled or unpopulated.
 - a. GREEN: Operating at PCI-Express Gen1 speed
 - b. YELLOW: Operating at PCI-Express Gen2 speed
 - c. RED: Operating at PCI-Express Gen3 speed
15. PCI-Express LED for PE3. The LED will remain off when this PCI-E slot is disabled or unpopulated.
 - a. GREEN: Operating at PCI-Express Gen1 speed
 - b. YELLOW: Operating at PCI-Express Gen2 speed
 - c. RED: Operating at PCI-Express Gen3 speed
16. PCI-Express LED for PE4. The LED will remain off when this PCI-E slot is disabled or unpopulated.
 - a. GREEN: Operating at PCI-Express Gen1 speed
 - b. YELLOW: Operating at PCI-Express Gen2 speed
 - c. RED: Operating at PCI-Express Gen3 speed
17. PCI-Express LED for PE5. The LED will remain off when this PCI-E slot is disabled or unpopulated.
 - a. GREEN: Operating at PCI-Express Gen1 speed
 - b. YELLOW: Operating at PCI-Express Gen2 speed
 - c. RED: Operating at PCI-Express Gen3 speed

Intel X299 FTW-K Motherboard Component Legend

The EVGA X299 FTW-K Motherboard with the Intel X299 and PCH Chipset.
Figure 1 shows the motherboard and Figure 2 shows the back panel connectors


FIGURE 1. X299 FTW-K Motherboard Layout




Component Legend		
1. CPU Socket 2066	12. M.2 Socket 3 Key-M 110mm	23. RGB Backlit Component Covers
2. Intel X299 PCH (Southbridge)	13. M.2 Socket 3 Key-M 80mm	24. Front Panel Audio Connector
3. CPU Fan Headers (1 amp PWM)	14. PCI-E Slot 16x/8x	25. Front Panel Connectors
4. Fan Headers (1 amp DC/PWM)	15. PCI-E Slot 4x	26. S/PDIF Out
5. RGB LED Controller Header	16. PCI-E Slot 1x	27. Removable BIOS Chip
6. DDR4 Memory DIMM Slots 1-8	17. Power Button	28. BIOS Selector Switch
7. 24-pin ATX power connector	18. Reset Button	29. CMOS Battery
8. 8 pin EPS Connector	19. Debug LED / CPU Temp Monitor	30. PC Speaker
9. Supplemental PCI-E 6 pin power	20. USB 3.0 Headers	31. Rear Panel Connectors (Figure 2)
10. Intel SATA 6G RAID Ports	21. USB 2.0 Headers	
11. U.2 (SFF-8639) Port	22. VROC Header	

**For a FULL description of the above legend, please see Page 15.

Figure 2. Chassis Rear Panel Connectors



I/O Hub		
1. USB 3.0	4. BIOS/CMOS Reset	7. M.2 Key-E Vertical Header
2. USB 3.1 type A	5. Killer E-2500 NIC	8. Optical Out
3. USB 3.1 type C	6. Intel I 219 NIC	9. Analog Audio Jacks



Speed/Link LED Status	Description
Orange	1000 Mbps data rate
Green	100 Mbps data rate
Off	10 Mbps data rate

Activity LED Status	Description
Off	No Data Transmission
Blinking (Green)	Data Transmission

3.5mm Audio Jack Legend				
Analog Audio Port Breakdown	2/2.1 Channel	4.0/4.1 Channel	5.1 (6 Channel)	7.1 (8 Channel)
Blue	Line in	Line In	Line in	* Rear Speaker Out
Green	Front Speaker Out/ Front Speaker + Sub	Front Speaker Out/ Front Speaker + Sub	Front Speaker Out	Front Speaker Out
Pink	Mic In	Mic In	Mic In	Mic In
Black		Side Speaker Out	Side Speaker Out	Side Speaker Out
Orange			Center/Sub Out	Center/Sub Out

* 7.1 output is enabled via Realtek Software from within Windows. (Page 95)

Component Legend Descriptions

1. CPU Socket 2066
This is the interface for the Central Processing Unit (CPU), and supports Core i7 models compatible with the Intel 2066 Socket Skylake-X and Kaby Lake-X architecture.
2. Intel X299 PCH (Southbridge)
The Platform Controller Hub (PCH) handles the role that was previously held by the South Bridge. The PCH has 4 PCI-E Gen 3 lanes and allocates bandwidth to smaller PCI-E slots, M.2 Key-E, USB, audio, etc. In simplified terms, the PCH works as a hub for peripherals that are less bandwidth-intensive.
3. CPU Fan Headers (PWM)
4-pin fan headers that control the fan speed based on a configurable curve or static percentage. PWM (Pulse-Width Modulation) works by pulsing power to the fan at a constant rate and sending the RPM signal to the fan's controller via a Sense cable, rather than adjusting fan speed by increasing and decreasing voltage. This method is preferable because it eliminates voltage-based fan stall points. Please see Page 87 for more in-depth PWM breakdown and PWM controls within BIOS/UEFI.
4. Fan Headers DC/PWM
These ports are simply toggled between DC and PWM to be compatible with 3 pin and 4 pin fans. The fans will be set to a static percentage set in BIOS.
5. RGB Header
The RGB header is a 4 pin header that allows a software based control within windows for RGB devices via ELEET-X. Please see Page 42 for control specifics.
6. DDR4 Memory Slots
The memory slots support up to four 288-pin DDR4 DIMMs in Quad-Channel mode with Skylake-X processors; and supports up to two 288-pin DDR4 DIMMs in Dual-Channel mode with Kaby Lake-X processors.

Skylake-X processors are certified for Quad-Channel mode, and will be enabled only upon using four sticks of supported memory, according to the installation guide on Page 28. Skylake-X supports up to 128GB (8x16GB) up to 4000MHz+, 32GB modules are *NOT* supported on this platform. Some

Dual-Channel kits *may* work; however Skylake-X is certified for Quad-Channel operation, not Dual-Channel. Dual-Channel configurations will substantially reduce the potential memory bandwidth of the Skylake-X processor.

Kaby Lake-X processors are certified for Dual-Channel mode, and will be enabled only upon using two sticks of supported memory, according to the installation guide on Page 28. Kaby Lake-X supports up to 64GB (4x16GB) up to 4133MHz+. 32GB modules are *NOT* supported on this platform. Some Quad-Channel kits *may* work, however Kaby Lake-X is not certified for Quad-Channel operation, only Dual-Channel, and will gain no benefit from Quad-Channel over Dual-Channel.

Using an odd number of DIMMs (1,3,5,7) will lower the board to Single-Channel mode, which may significantly lower performance depending on the application.

The speeds listed above cannot be guaranteed as Intel® only certifies the speed of the memory controller up to 2666/2400MHz for Kaby Lake-X and Skylake-X platforms respectively, and all speeds above the speeds certified by Intel® require overclocking.

7. 24-pin ATX power connector

The main power for the motherboard is located on the right side of the board and perpendicular to the PCB; this is also described as a “Vertical” connector (See Page 38 for more specifics to the connector itself, and associated wiring/pinouts). The 24-pin connector IS directional and the connector needs the tab on the socket to line up with the release clip located on the 24-pin connector from the power supply. This connector pulls the bulk of the power for all components; other connectors, such as EPS, PCI-E (video card AND motherboard sides), have been added to reduce the load and increase longevity due to wiring and trace limitations.

8. 8-pin EPS Connector

The EPS is dedicated power for the CPU (See Page 39 for more specifics to the connector itself, and associated wiring/pinouts). Carefully choose the correct power cable by consulting with the installation manual for your power supply. This connector is designed only to work with an EPS or CPU cable. System builders may make the mistake of plugging in a PCI-E 8-pin or 6+2-pin connector, which will prevent the board from POSTing and possibly damage the board; although the cables appear similar, they are wired differently and attaching a PCI-E cable to an EPS connector may cause damage to the motherboard.

Alternatively, if no power cable is connected or detected, the system will not POST and will hang at POST code “C.”

9. Supplemental PCI-E 6-pin Power Connector

There is a 6-pin PCI-E connector at the bottom of the motherboard (See page 38 for more specifics to the connector itself, and associated wiring/pinouts). This connector provides dedicated power to the PCI-E x16 slots, augmenting the power provided by the 24-pin and the GPU directly.

This is optional for a single card solution, and is recommended for SLI, CFX, and dual processor video cards.

10. Intel SATA 3/6G Ports

The Intel X299 PCH has a 8-port SATA 3/6G controller (See Page 46 for specifics on the connectors). This controller is backwards-compatible with SATA and SATA 2 devices, and supports SSDs, HDDs and various types of optical devices (CDROM, DVDROM, BD-ROM, etc). The controller also supports NCQ, TRIM, hot swap capability (provided the proper HDD/SSD bays/racks are installed), and RAID levels 0/1/5/10.

11. U.2 Port (SFF-8639)

U.2, originally known as SFF-8639, is a high bandwidth connection specifically engineered for next generation SSD's. U.2 brings PCI-E x4 (Gen3) NVMe performance to a 2.5" SSD form factor and provides a solution to potential heating problems that may be present in some M.2 solutions.

12. M.2 Socket 3 Key-M 110mm

M.2 is a SSD standard, which uses up to four PCI-E lanes and utilizes Gen3 speeds. Most popularly paired with NVMe SSDs, this standard offers substantially faster transfer speeds and seek time than SATA interface standards. All M.2 devices are designed to connect via a card-bus style connector and be bolted into place and powered by the connector, rather than by a dedicated data cable and power cable. This socket will support Key-M devices of 110mm, 80mm, 60mm, and 42mm length.

This connector can utilize either a PCI-E/NVMe based M.2 SSD or SATA M.2.

13. M.2 Socket 3 Key-M 80mm

M.2 is a SSD standard, which uses up to four PCI-E lanes and utilizes Gen3 speeds. Most popularly paired with NVMe SSDs, this standard offers substantially faster transfer speeds and seek time than SATA interface standards. All M.2 devices are designed to connect via a card-bus style

connector and be bolted into place and powered by the connector, rather than by a dedicated data cable and power cable.
 This socket will support Key-M devices of 80mm, 60mm, and 42mm length.
 This connector can utilize either a PCI-E/NVMe based M.2 SSD, SATA M.2, or Intel Optane devices.

14. PCI-E Slot x16/x8*

PCI-E x16/x8 slots are primarily for video cards. These full-length slots will provide 8 or 16 lanes of bandwidth to a full-size card, and are backwards-compatible with x8, x4, and x1-length cards.
 Skyake-X Socket 2066 processors have 44 or 28 PCI-E lanes available for routing, whereas Kaby Lake-X has 16 PCI-E lanes.

15. PCI-E Slot x4*

PCI-E x4 slot uses up to 4 Gen 3 lanes from the PCH. This slot is typically used for sound cards, WiFi, USB, or other peripheral cards.
 Because this slot uses PCH bandwidth, this will have *NO EFFECT* on the bandwidth or throughput of the x16 slots used for SLI.

16. PCI-E Slot x1*

PCI-E 1x is the smallest form-factor PCI-E card slot. They are all one lane and are used for low-bandwidth products. This slot uses 1 lane from the PCH.

17. Power Button

This is an onboard power button, and may be used in place of, or in conjunction with, a front panel power button wired to the board.
 Benching systems, or test benches before final assembly, are best served by using the onboard power because it removes the need to wire a Power/Reset button or cross posts with a screwdriver, which is a semi-common practice.
 This button provides a safer and easier option than jumpering the Power posts.

18. Reset Button

This is an onboard system reset button, and may be used in place of, or in conjunction with, a front panel system reset button wired to the board.
 Benching systems, or test benches before final assembly, are best served by using the onboard power because it removes the need to wire a Power/Reset button or cross posts with a screwdriver, which is a semi-common practice.
 This button provides a safer and easier option than jumpering the Power posts.

19. Debug LED / CPU Temp

This is a two-digit POST code reader, displaying in hexadecimal, which means the characters available (when working as intended) are 0-9, A-F and has a cap of 255 characters. The POST codes are listed in the troubleshooting section on

Page 142. After the system boots, it will display the temperature in Celsius. This temperature is specifically for the CPU socket, which will typically read slightly higher than a given CPU core. To read this temp in Fahrenheit, take the value in Celsius, multiply by 9/5 (or 1.8) and add 32.

20. USB 3.0 Headers

The USB3.0 headers are used to connect additional USB interface plugs to the motherboard; these headers are most often used to connect the motherboard to the chassis to enable the USB3.0 ports on the chassis. These will function the same as the USB3 ports found on the motherboard's hardwired I/O hub, but these can be used to attach to front panel USB, auxiliary ports that mount in the card slots, and also some devices that directly connect to the header.

USB 3.0 standard is 900ma @ 5V for unpowered devices. If your USB device requires more power than this, it is recommended to attach a powered USB Hub.

USB 3.1 Type-A (found on the I/O Hub) shares the power limit of USB 3.0 at 900ma @ 5V. Whereas USB 3.1 Type-C (also found on the IO Hub) has a power limit of 3000ma (3A) @ 5V.

21. USB 2.0 Headers

The USB2.0 header is used to connect additional USB interface plugs to the motherboard; these headers are most often used to connect the motherboard to the chassis to enable the USB2.0 ports on the chassis. These will function the same as the USB2 ports found on the motherboard's hardwired I/O hub, but these can be used to attach to front panel USB, auxiliary ports that mount in the card slots, and also some devices that directly connect to the header.

USB 2.0 standard is 500ma @ 5V per port (header total is 1000ma) for unpowered devices. If your USB device requires more power than this, it is recommended to attach a powered USB Hub.

22. VROC Header

VROC stands for Virtual RAID On CPU, the VROC headers works in conjunction with the upcoming VROC cards. VROC cards are 4 device M.2 Key-M cards for PCI-E that allow RAID functions on the card. The Header is for an Intel hardware key that will unlock advanced RAID functions, which in VROC's case is anything other than RAID0.

Important note, as of the time this manual was written, VROC will work with many SSD's but is only bootable with Intel SSD's. Also VROC is only compatible with Skylake-X CPUs

23. RGB Backlit Panels

These panels sit on top of other components, such as heatsink, and provide RGB LED lighting. These panels connect to the RGB Headers (component #5) and provide stylized backlighting which can be controlled in EVGA ELEET-X.

These RGB Headers are **ONLY** to be used for the included/preattached RGB covers, not a 3rd party RGB device.

24. Front Panel Audio Connector

This is a motherboard header, which is used to plug in the audio cable originating from most PC chassis to allow audio to be recorded from or played through the audio connectors on the chassis. This header has a connector that looks similar to USB2 and will use the standard “HD Audio” jack. Some cases may have two headers on one cable strand: one labeled HD Audio, and one labeled AC’97 – this header is **not** compatible with AC’97.

25. Front Panel Connectors

The Front panel connectors are the four main chassis connections. These include the Power Switch, Power LED, Reset Switch, and HDD LED. The Power and Reset switches are both designed to use “Momentary Switches,” rather than “Latching Switches,” which means the connection between the two posts needs to be made just briefly for it to work, as opposed to being held in place. This is why the Power and Reset switches can be triggered with a screw driver by simultaneously touching the + and - posts.

Power LED will power on with the system, indicating the system is on and can blink with CPU activity.

HDD LED will blink during access to the SATA ports, U.2 port, and/or the M.2 Key-M SSD’s.

26. S/PDIF Out

S/PDIF is used for audio pass-through, which is used on some older video cards, video capture cards, and some older generation devices. This port is not widely used now, but is here in case a legacy item requires the connector. Most of the audio data is pushed through the bus now, and does not require an external cable. NOTE: This is a S/PDIF **Out** connector, and cannot operate to record data from an external device.

27. Removable BIOS Chip

The ability to remove and replace the BIOS chip without requiring special tools is a key feature on many of our boards, by adding a level of protection against BIOS failures, bad BIOS flashes, BIOS corruption, etc. This feature also allows EVGA to work with end-users if something happens to render the BIOS chip unusable or a BIOS update is needed for CPU compatibility; rather than

replacing the whole board, Customer Support can send a pre-flashed BIOS chip to get you up and running without requiring the end-user to send in the motherboard. Please see the walkthrough on Page 132 for further information.

28. BIOS Selector Switch

This switch toggles between physical BIOS chips. This board has a BIOS chip as a permanent fixture, and the other is a replaceable chip listed above. Each chip holds profiles and all settings that have been saved to the BIOS. This allows you to swap between two physically different BIOS chips. This also makes it substantially easier to hotflash a BIOS if needed. If instructions are needed for hotflashing a BIOS, please EVGA Customer Service (Page 139 for contact info).

29. CMOS Battery

The CMOS battery backup provides uninterruptable power to the BIOS/UEFI to keep all of the settings; otherwise, each boot would behave like you just reset the BIOS. These batteries typically last several years and rarely need to be replaced.

30. PC Speaker

This is a small mono low-fidelity speaker permanently attached to the motherboard used mainly for debugging purposes. A POST beep may indicate a successful POST, various tones for USB initialization, and other beeps to indicate an issue during the post process. Please see Page 140 for more details.

31. Rear Panel Connectors (Figure 2)

This is the section referred to as the I/O Hub. This panel contains the hardwired USB, Sound, and Ethernet connections. Please see Page 14 for a component level breakdown.

CMOS Reset Button (On IO Panel)

This button has two main uses: the first is standard practice to clear BIOS and power on before updating the BIOS, and the second is standard practice when troubleshooting instances when the motherboard fails to POST, such as after upgrading RAM or CPU, installing new hardware, a failed overclock, etc. This button provides a much faster means of resetting than the previous method of removing power from the board, removing the CMOS battery, and discharging power to the board. In rare occasions the older method can help; pressing the clear CMOS button will normally allow you and your system back into the default BIOS.

M.2 Socket 1 Key-E 32mm Vertical Adapter (On IO Panel)

M.2 Key-E is largely used for WiFi and Bluetooth cards. Key-E and Key-M connectors are different, meaning that devices are not interchangeable between sockets.


This variant is also on a vertical mount, which has an adapter that is slid into card-bus slot and bolted to the motherboard, then the M.2 Key-E mounts to the vertical adapter.

* There are two numeric references for PCI-Express: one is mechanical, which is the actual slot-length footprint, and the second is electrical, which is a reference of how many PCI-E lanes are routed to the slot.

As PCI Express is designed to be a universal architecture, you can install x1 cards, such as sound cards or USB controllers into an x16 slot. Many types of cards can use different amounts of PCI-E lanes, while some applications use only certain parts of a card, such as compute apps that allow a card to run off of a single PCI-E lane. This is why there are x16 mechanical slots with x1 electrical PCI-E lane. Using the entire length of a PCI-E slot is not necessary, nor does it cause an adverse effect to use a shorter form-factor bus card in a slot that physically can hold a larger form-factor bus card.

Card Slots

The X299 FTW-K features four x16 PCI-E slots, one x4 PCI-E slot, one x1 PCI-E slot, one Socket 3 Key-M M.2 110mm (backwards compatible with Key-M 80mm, 60mm, and 42mm), and one Socket 3 Key-M M.2 80mm (backwards compatible with Key-M 60mm and 42mm).



PCI-E Slot Breakdown

PCI-E Lane Distribution (44 Lane SKX CPU's)

- ❑ PE1 – x4 (Gen3, x4 lanes from PCH)
- ❑ PE2 – x16 (Gen3, x16 lanes from CPU, x8 shared with PE3)
- ❑ PE3 – x16 (Gen3, x8 lanes from CPU, shares 8 of PE2's 16 lanes)
- ❑ PE4 – x16 (Gen3, x8 lanes from CPU, shares 8 of PE5's 16 lanes)
- ❑ PE5 – x16 (Gen3, x16 lanes from CPU, x8 shared with PE4)
- ❑ PE6 – x1 (Gen3, x1 lane from PCH)

PCI-E Lane Distribution (28 Lane SKX CPU's)

- ❑ PE1 – x4 (Gen3, x4 lanes from PCH)
- ❑ PE2 – x16 (Gen3, x16 lanes from CPU, x8 shared with PE3)
- ❑ PE3 – x16 (Gen3, x8 lanes from CPU, shares 8 of PE2's 16 lanes)
- ❑ PE4 – x16 (Not functional with a 28 lane processor.)
- ❑ PE5 – x16 (Gen3, x8 lanes from CPU)
- ❑ PE6 – x1 (Gen3, x1 lane from PCH)

PCI-E Lane Distribution (16 Lane KBX CPU's)

- ❑ PE1 – x4 (Gen3, x4 lanes from PCH)
- ❑ PE2 – x16 (Gen3, x8 lanes from CPU)
- ❑ PE3 – x16 (Not functional with a 16 lane processor.)
- ❑ PE4 – x16 (Not functional with a 16 lane processor.)
- ❑ PE5 – x16 (Gen3, x8 lanes from CPU, disabled if 110mm M.2 is used)
- ❑ PE6 – x1 (Gen3, x1 lane from PCH)

M.2 and U.2 Slot Breakdown

PCI-E Lane Distribution (44 Lane SKX CPU's)

- ❑ U.2 1 – Gen3, x4 lanes from CPU (No Shared Lanes)
- ❑ U.2 2 – Gen3, x4 lanes from CPU (No Shared Lanes)
- ❑ M.2 Key-M (110mm) – x4 CPU lanes (No Shared Lanes)
- ❑ M.2 Key-M (80mm) – x4 PCH lanes (No Shared Lanes) (*)
- ❑ M.2 Key-E (32mm) – x1 PCH lane (No Shared Lanes)

PCI-E Lane Distribution (28 Lane SKX CPU's)

- ❑ U.2 1 – Not functional with a 28/16 lane processor.
- ❑ U.2 2 – Not functional with a 28/16 lane processor.
- ❑ M.2 Key-M (110mm) – x4 CPU lanes (No Shared Lanes)
- ❑ M.2 Key-M (80mm) – x4 PCH lanes (No Shared Lanes) (*)
- ❑ M.2 Key-E (32mm) – x1 PCH lane (No Shared Lanes)

PCI-E Lane Distribution (16 Lane KBX CPU's)

- ❑ U.2 1 – Not functional with a 28/16 lane processor.
- ❑ U.2 2 – Not functional with a 28/16 lane processor.
- ❑ M.2 Key-M (110mm) – x4 CPU lanes (Gen3, x4 shared with PE5)
- ❑ M.2 Key-M (80mm) – x4 PCH lanes (*)
- ❑ M.2 Key-E (32mm) – x1 PCH lane

- All M.2 Key-M slots on this board support PCI-E, NVMe, and SATA M.2 standards.
 - (*) If a SATA M.2 drive is used in the 80mm M.2 the SATA 0/1 will be disabled.
- This motherboard does NOT have any lane replication via PLX; all lanes are native and derived from CPU or PCH. This also allows for improved backwards compatibility for Gen 2 devices.



Installing the CPU

Note: Use extreme caution when working with the CPU to not damage any pins in the CPU socket on the motherboard!

Note: After removing the CPU socket cover, it is recommended to store it in case you ever need to transport your motherboard. If you ever remove the CPU, it is highly recommended to reinstall the socket cover.


Use the following procedure to install the CPU onto the motherboard:

1. Remove the plastic protective socket cover by pulling it straight up. Be sure not to damage any of the pins inside the socket.
2. Unhook the left socket lever by pushing down and away from the socket.
3. Unhook the right socket lever, this will release the load plate
4. Pull the socket lever back and the load plate will lift.
5. Open the load plate and make sure not to damage any of the pins inside the socket.



6. Align the notches on the CPU to the notches in the socket.
7. Lower the processor straight down into the socket.

Note: Make sure the CPU is fully seated and level in the socket before lowering the load plate.



8. Lower the load plate so it is resting on the CPU.
9. Carefully lock the lever back into place.



Installing the CPU Cooling Device

There are many different cooling devices that can be used with this motherboard. Follow the instructions that come with your cooling assembly.

Installing System Memory

Your X299 FTW-K has (8) 288-pin slots for DDR4 memory. These slots support 4GB, 8GB and 16GB DDR4 DIMMs. There must be at least one memory slot populated for the board to boot and operate.

The Intel X299 chipset supports Quad-Channel DDR4 memory, a maximum of 128GB and up to 4000MHz+ for Skylake-X, and 4133MHz+ for Kaby Lake-X (OC)+. It is recommended to always use a 4 DIMM Quad Channel kit for Skylake-X, and 2 DIMM Dual Channel kit for Kaby Lake-X; in the event you cannot please fill the memory slots in the following order: 1,3,5,7,2,4,6,8 (5,7,6,8 for KBX). See chart to the right: Use the following procedure to install memory DIMMs. Note that there is an off-center gap near the center of the DIMM slots. This matches the gap on a DDR4 DIMM to ensure the memory is installed properly, and to prevent the incorrect installation of memory.

1. Unlock a DIMM slot by pressing the top-side module clips outward.
2. Align the memory module to the DIMM slot, and insert the module perpendicular to the DIMM slot, pressing straight down to seat the module. The plastic clips at top side of the DIMM slot automatically lock the DIMM into the connector.



Skylake-X RAM Slot Fill Order								
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
1 DIMM	X							
2 DIMM	X		X					
3 DIMM	X		X		X			
4 DIMM	X		X		X		X	
5 DIMM	X	X	X		X		X	
6 DIMM	X	X	X	X	X		X	
7 DIMM	X	X	X	X	X	X	X	
8 DIMM	X	X	X	X	X	X	X	X

Kaby Lake-X RAM Slot Fill Order								
	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
1 DIMM	N/A	N/A	N/A	N/A	X			
2 DIMM	N/A	N/A	N/A	N/A	X		X	
3 DIMM	N/A	N/A	N/A	N/A	X	X	X	
4 DIMM	N/A	N/A	N/A	N/A	X	X	X	X
5 DIMM	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
6 DIMM	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
7 DIMM	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
8 DIMM	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Note: The memory controller on most Skylake-X and Kaby Lake-X CPUs runs at a default frequency of 2400/2666MHz. Achieving memory speeds above 2666MHz+ may require manual setting of the memory timings, frequency and voltages and/or overclocking of the CPU.

Refer to the memory manufacturer specifications for the recommended memory timings.

Installing the I/O Shield and I/O Cover

The motherboard kit comes with an I/O shield that is used to block internal components from dust and foreign objects, while also promoting correct airflow within the chassis.

Before installing the motherboard, install the I/O shield from the inside of the chassis. Press the I/O shield into place and make sure it fits securely.

The X299 FTW-K also includes an I/O cover. This I/O cover adds a unique appearance to the I/O area of the motherboard and is completely optional. If you wish to use the cover, please place it over the I/O area and install the chassis screws. These screws will secure the I/O cover to the motherboard.

Installing the Motherboard

Installing the motherboard into a system case depends on several factors: whether you are replacing an existing motherboard, whether you are building a new PC, and the type of chassis that will house your PC components. You must first determine if it would be easier to secure the motherboard to the chassis or if it would be easier to install other components prior to this step. It is normally easier to secure the motherboard first.

Note: Be sure that the CPU fan assembly has enough clearance for the system case covers to lock into place and for expansion cards. Also, make sure the CPU Fan assembly is aligned with the vents on the covers; this is to allow the airflow to properly exhaust from the chassis. The CPU Fan assembly orientation will depend on the system case being used.


Securing the Motherboard into a System Case

Most system cases require installation of standoffs into the chassis to allow the motherboard to be mounted to the chassis and prevent short circuits. If there are studs that do not align with a mounting hole on the motherboard, it is recommended that you remove that standoff to prevent the possibility of a short circuit; also ensure that all needed standoffs are fully tightened before attaching the motherboard to the chassis. Please review the installation manual included with your chassis for the proper installation of the motherboard standoffs.

1. Carefully place the motherboard onto the standoffs located inside the chassis.
2. Align the mounting holes with the standoffs.
3. Align the connectors to the I/O shield and/or I/O cover.
4. Ensure that the fan assembly is aligned with the chassis vents according to the fan assembly instruction.
5. Secure the motherboard with ten (10) screws (See next page for mount hole location). Ensure that each screw is lined up with and screwing into the corresponding standoff under the board. Double-check alignment to make sure nothing gets cross-threaded.

Tip: If you have difficulty with getting some of the screws fastened, especially near the I/O hub, first try to loosely fasten all other screws on the motherboard, but don't completely tighten the screws. This may help to hold the board in place, allowing you to thread and fasten the remaining screws. Once all screws are properly threaded, remember to go back and tighten the rest of the screws.

6. See the picture below for a zoomed-in view of a hole to place over a standoff, as well as the locations of standoff holes for the X299 FTW-K.




1. All safe locations to secure the board to a standoff are circled above in red.
2. Keep in mind that when the screws are installed, but not fully tightened, the motherboard should have 1-2mm of movement; this can help with mounting cards or tight-fits with other components.
3. Once the board is properly aligned, be sure to fully-tighten the board to the chassis before proceeding.

Installing M.2 devices

Securing a M.2 device to the motherboard requires a few extra steps compared to other current drive or slot-based connectors. M.2 devices used on this motherboard, Socket 3 (for SSDs) and Socket 1 (for WiFi/Bluetooth), will be installed and attached the same way. Below are images from an installation of a SSD on a Socket 3 Key-M.

1. Remove the screw that comes pre-attached to the M.2's retention standoff; this will be used to keep the device in place.




2. Insert the M.2 device in at a slight angle, at approximately 45 degrees to the board. This will allow the contacts (colloquially called “Gold Fingers”) to seat completely into the slot. If the device is fully seated, you should be able to release it and the device will rest at an angle of about 30 degrees on its own, as shown in the picture below.



3. Gently push the M.2 device down on the raised end. There will be some tension - this is normal - then use the screw you removed in Step 2 to secure the device.

Below, you can see that the contacts will be nearly invisible when the device is properly seated and the copper mounting semi-circle is partially visible around the screw.




Incorrect installation Example:

***NOTE* This is one of the most common examples of an incorrect installation of a M.2 device. Do not intentionally attempt this, or complete your installation with this example. Doing so could cause damage to the device or the M.2 port.**

Using the image below, notice how the gold fingers fail to fully seat in the M.2 slot. This often occurs if the drive was pushed into the slot from a nearly parallel starting position, rather than an angle, causing the M.2 device to not seat fully. As a result, this may cause a drive detection failure by the BIOS, the drive will be detected with the description in gibberish (e.g. characters in the name, such as @, #, \$, %, *, etc.) and/or notifications that the drive has corrupt data stored on it.

Moreover, if the device's connection looks like this, then screwing down the device will likely be a VERY tight fit, if even possible, which is another sign that the device is not seated properly.

In conclusion, if the install device looks similar to the image below, please remove and reseal it using the instructions above. **DO NOT POWER THE SYSTEM ON IF THE CONNECTOR RESEMBLES THIS PICTURE.**



For further M.2 setup instructions please see Page 50.

Tested CPU

CPU	Core Count	Frequency	PCI-E Lanes
Skylake-X Core Series			
Core™ i9-7980XE	18 Cores + HT	TBD	44
Core™ i9-7960X	16 Cores + HT	TBD	44
Core™ i9-7940X	14 Cores + HT	TBD	44
Core™ i9-7920X	12 Cores + HT	TBD	44
Core™ i9-7900X	10 Cores + HT	3.3 GHz \ 4.5 GHz (Boost)	44
Core™ i9-7820X	8 Cores + HT	3.6 GHz \ 4.5 GHz (Boost)	28
Core™ i9-7800X	6 Cores + HT	3.5 GHz \ 4.0 GHz (Boost)	28
Kaby Lake-X Core Series			
Core™ i7-7740K	4 Cores + HT	4.2 GHz \ 4.5 GHz (Boost)	16
Core™ i5-7640K	4 Cores	4.0 GHz \ 4.0 GHz+ (Boost)	16

Tested U.2

U.2 (SSD):			
Brand	Part Number	Size	Interface
Intel	750 SSDPE2MW400G4	400GB	U.2 NVME w\Cable
Intel	750 SSDPE2MW400G4	400GB	U.2 NVME w\Cable

Tested M.2 Key-M

M.2 Key M (SSD) :			
Brand	Part Number	Size	Interface
Samsung	MZ-HPU128T/004 128G	128GB	M.2
Samsung	MZ-HPV1280 SM951	128GB	M.2
Intel	SSDSCKHW120A4	120GB	M.2
Transcend	TS128GMTS800	128GB	M.2
Samsung	MZ-HPU128T/004	128GB	M.2
Kingston	SHPM2280P2H/240	240GB	M.2 + PCIE Adapter Card
Samsung	MZ-VKV512 NVM Express	512GB	M.2
Intel	SSDSCKKW120H6	120GB	M.2
Intel	SSDSCKKW240H6	240GB	M.2
Intel	Intel SSD 6000P SERIES SSDPEKKW256G7	256GB	M.2 + PCIE Adapter Card
Intel	Intel SSD 6000P SERIES SSDPEKKW512G7	512GB	M.2 + PCIE Adapter Card

Tested M.2 Key-E

M.2 Key E (WiFi):		
Brand	Part Number	Wifi Support
Intel	7265NGW	BT 4.0, 802.11ac
Intel	8260NGW	BT 4.1, 802.11sc
Azure Wave	AW-NB165NF	BT 4.0, WLAN B, G, N

Connecting Cables


Important caveat: the images below do **NOT necessarily represent the physical orientation of the specific header on *your* motherboard. Rather, these are physical representations designed to give you a basic physical footprint of the component AND the cable pinouts for the ports.**

The locations of these ports can be found in the Component Legend on Page 12-13.

This section takes you through all the necessary connections on the motherboard. This will include:

- ❑ Power Connections
 - 24pin ATX power (**PW1**)
 - EPS 8-pin 12V power
- ❑ Internal Headers
 - Front Panel connectors (Power/Reset/LEDs)
 - Fan Headers (PWM for CPU and DC for case fans)
 - USB 2.0 Header
 - USB 3.0 Header
 - Audio Header
 - SATA
 - U.2
- ❑ Rear I/O Panel


24pin ATX Power (ATX_PWR_24P)

24-pin Main Power			
			
			
1 +3.3V	7 GND	13 3.3V	19 GND
2 +3.3V	8 PWROK	14 -12V	20 RSVD
3 GND	9 +5V_AUX	15 GND	21 +5V
4 +5V	10 +12V	16 PS_ON	22 +5V
5 GND	11 +12V	17 GND	23 +5V
6 +5V	12 +3.3V	18 GND	24 GND

ATX_PWR_24p is the main power supply connector located along the lower-right edge of the board. Make sure that the power supply cable and pins are properly aligned with the connector on the motherboard. Firmly plug the power supply cable into the connector and make sure it is secure.


The 24-pin Power Connector may be standard or right-angled depending on your motherboard model.

6-pin PCI-E

ATX_PWR_6P	
	
1 +12V	
2 +12V	
3 +12V	
4 GND	
5 GND	
6 GND	


The 6-pin PCI-E connector present on the motherboard provides additional power to the PCI-E slots, rather than pulling it all from the 24-pin main power. It is advised to plug in this connector when using SLI, especially with higher-end graphic cards. It is strongly recommended to plug in this connector for 3-way and 4-way SLI, regardless of GPU.

EPS 8-pin 12V Power (PWR 8P)



EPS PWR 8P, the 8-pin ATX 12V power connection(s), is used to provide power to the CPU. Align the pins to the connector and press firmly until seated. The secondary EPS, if present, is optional for improved overclocking. Please remember to make sure that the tab on the EPS socket is aligned with the release clip on the cable. **NOTE:** If the tab and release clip are on opposite sides, yet the power connector fits, then you are using a PCI-E 8-pin cable, which **WILL** damage the board if powered on. Please review installation instructions from your power supply manufacturer to verify which connectors may be used for the CPU power.

Front Panel Header



The front panel header on this motherboard is used to connect the following four cables:

❑ **PWRLED**

Attach the front panel power LED cable to these two pins of the connector. The Power LED indicates the system's status. When the system is powered on, the LED will be on.

❑ **PWRSW**

Attach the power button cable from the case to these two pins. Pressing the power button on the front panel turns the system on and off rather than using the onboard button.

❑ **HD_LED**


Attach the hard disk drive indicator LED cable to these two pins. The HDD indicator LED indicates the activity status of the hard disks.

❑ **RESET**

Attach the Reset switch cable from the front panel of the case to these two pins.


Note: Some system cases may not have all four cables. Be sure to match the name on the connectors to the corresponding pins.

Fan Header




This motherboard line only has 4-pin fan headers, which are backwards compatible with 3-pin fan connectors. Fans may be controlled by PWM or DC controls. The headers have an absolute safe power limit of 1 Amp @ 12 Volts (12 Watts). These headers are for your CPU heatsink and chassis cooling fans.

RGB LED Header



The RGB LED header supports included EVGA RGB covers, NOT to be used with other RGB devices. The LEDs can be controlled from within ELEET-X.


VROC Header



This is the hardware key for enabling advanced RAID features (anything above RAID 0) for “Virtual RAID On CPU” based M.2 PCI-E cards.

USB Headers

The motherboard contains 1x 19-pin internal header connectors onboard that can be used to connect an optional external bracket containing up to four (2) USB 3.0 ports.




- ❑ Secure the bracket to either the front or rear panel of your chassis (not all chassis are equipped with the front panel option).
- ❑ Connect the end of the cable to the USB 3.0 headers on the motherboard.

The motherboard contains 2x 9-pin internal header connectors onboard that can be used to connect an optional external bracket containing up to four (4) USB 2.0 ports.

USB 2.0 Header






Pin	Signal
1	5V_Dual
3	D-
5	D+
7	GND
9	Empty
2	5V_Dual
4	D-
6	D+
8	GND
10	No Connect


- ❑ Secure the bracket to either the front or rear panel of your chassis (not all chassis are equipped with the front panel option).
- ❑ Connect the end(s) of the cables to the USB 2.0 headers on the motherboard.

Front Panel Audio Header



Front panel audio supports HD Audio for stereo/gaming headphones or 2.1 speakers, and a Mic.


S/PDIF Header



The S/PDIF (Sony/Philips Digital Interface Format) is used for specific Audio-Out applications, typically for using an audio pass-through device.

Drive Headers (SATA/ U.2)

SATA3/6G is the current standard for HDD/SSD/Optical interface. These cables are the data interconnect for the motherboard. Your HDD/SSD/Optical interface will still require a separate power connection from your power supply.



SATA ports on this platform natively support full AHCI and RAID functions. AHCI is enabled by default, but the controller can be put into RAID mode in the BIOS. RAID mode supports RAID levels 0, 1, 5, and 10 through the Intel controller. RAID-ready ports also have full AHCI functionality and are boxed above in GREEN.

U.2 is a new storage standard that has the advantage of the performance of a M.2 SSD and the ease of installation of a 2.5inch SSD. U.2 does not share resources with other peripherals or slots, U.2 has a dedicated 4 PCI-E Gen3 lanes from the PCH, and can be enabled or disabled in BIOS.

See Page 52 for RAID levels supported and explanations for how they work.

Onboard Buttons

These onboard buttons include RESET, POWER and Clear CMOS. These functions allow you to easily reset the system, turn on/off the system, or clear the CMOS.

Clear CMOS Button

The motherboard uses CMOS RAM to store set parameters. The CMOS may be cleared by pressing the Clear CMOS button on the motherboard PCB or on the external I/O Panel.



External Clear CMOS Button

RESET and POWER Button

These onboard buttons allow you to easily turn on/off the system. These buttons allow for easy debugging and testing of the system during troubleshooting situations.

The POWER button has an integrated LED, indicating the system's status. When the system is powered on, the LED remains a solid red.

The RESET button has an integrated LED, indicating the activity status of the hard disk drives and will flicker accordingly.




↑
Reset
Button

↑
Power
Button

First Boot

When you power the system on for the first time (or after a BIOS update/reset) it may take a little longer than expected, and follow with a pause and message on the screen reading “BIOS Checksum error, Press F2 to continue or F12 to enter the BIOS.” This is completely normal.



Entering Setup...
CMOS checksum error,
Press F1 Skip, F12 Enter Setup

Depending on when the motherboard was purchased, compared to when it was released, there may be a BIOS update for it. If you have internet access on another PC and an USB Thumb Drive, it is highly recommended to download the current BIOS and flash the newer BIOS before attempting anything else.

If there is an updated BIOS, download and extract it onto a thumb drive formatted for FAT32, go to the Boot section, under “BIOS Firmware Update” select “Select BIOS File” and it will allow you to navigate to where your ROM file was stored, and press Enter to update.

****IMPORTANT NOTE. DO NOT CUT POWER OR REBOOT THE SYSTEM DURING A BIOS UPDATE; YOU WILL LIKELY RENDER THAT SPECIFIC BIOS UNBOOTABLE****

If you did update the BIOS, you will be greeted with the same screen as before stating that there is a Checksum error. Please press “Delete” to go into the BIOS/UEFI.

The X299 chipset is designed for UEFI mode and Windows 10 natively. However, you can run in Legacy mode, if need be. Typically, there is very little setup needed on these boards.

To setup the Memory, use the arrow keys or your mouse to select the “Memory” setting, select the “Memory Profiles” pulldown and select XMP Profile 1. This will automatically set some basic memory information, such as speed, latency and voltage. You may set the memory speed manually, if you’d prefer, but the XMP will generally get the memory running at the memory manufacturer’s specification with little to no effort.

Next, click “Boot” from the menu list at the top. “Boot Option #1” should show the device that you intend to install your operating system. If the device does not show up in this menu, and is a standard SSD/HDD plugged into a SATA port, proceed down to “UEFI Hard Disk Drive BBS Priorities” at the

bottom. In this menu, the top item will be “1st Boot” and will have a pulldown menu on the right. Click on the pulldown menu and select the intended drive; this will make the drive appear on the previous menu. If the drive doesn’t show at this point, please check the troubleshooting section on Page 135.

If you plan on using the M.2/U.2/SATA Express as a boot device, click on or navigate to the “Advanced” menu, select “Onboard Device Configuration” and enable the desired port.

Once this is done, press F10 to save and exit, plug in your operating system installation medium (likely a thumb drive) and Windows 10 should be able to boot to M.2 / U.2 / SATA Express without issue.

Note Some device manufacturers require specific drivers for HDDs or SSDs (such as M.2) before Windows can detect the drive for installation. Please make sure to consult the manufacturer’s instructions for your HDD or SSD before attempting to install Windows to determine if additional drivers are needed.

M.2 SSD, PCI-E SSD, and NVMe SSD Installation steps

M.2 is a card bus that can use multiple connector types to connect many types of devices, such as WiFi or SSDs, in a very small and power efficient package that is also very fast. M.2 devices can be connected to a M.2 card slot on the motherboard or through PCI-E on a M.2 to PCI-E adapter. There are also PCI-E native drives of this nature, such as the Intel 750. NVMe stands for Non-Volatile Memory Express, and is a high-performance standard for M.2/PCI-E SSDs. Most NVMe drives will require their own device drivers during the Operating System installation to be used as a boot drive, whether in M.2 or PCI-E form.

M.2/PCI-E share resources with other components on the motherboard, which may limit some hardware combinations. **PLEASE SEE CHART ON PAGE 25 FOR A DETAILED BREAKDOWN OF SHARED RESOURCES.**

M.2 Physical Installation

1. Please see Page 32 for M.2 physical installation instructions.

PCI-E Physical Installation

1. Install the SSD into any available PCI-E slot with at least x4 Lanes available (Gen3 Lanes are preferred for better speed, but Gen2 Lanes will still outperform the top end of SATA SSDs).
 - a. Please refer to Page 25 under PCI-E Lane Breakdown for lane specifics, as the PCI-E slot must be at least x4 electrically.
2. Attach the SSD's mounting bracket to the back of the case.
3. Native PCI-E drives will not normally require a separate power or data cable attached; all power and data transfers are done through the PCI-E slot.

BIOS Setup and Windows Installation for M.2 and PCI-E NVMe SSDs

1. Remember, NVMe is a new standard and older operating systems do not have native support. Many NVMe drives require certain steps to make the drive bootable, even with current operating systems.
PLEASE FULLY READ THE INSTRUCTIONS THAT COME WITH YOUR M.2 or PCI-E NVMe SSD BEFORE INSTALLATION.
2. After reviewing your SSD's instructions and its respective Physical installation instructions above, power on the PC and enter the BIOS/UEFI by pressing the F2 key repeatedly.

3. Once in BIOS/UEFI, navigate to the “BOOT” section. Then go down to the “CSM Configuration” heading and press enter, or click on it with your mouse.
 - a. For Windows 10: Set “Launch Storage OpROM Policy” to “UEFI”. Then set “Launch CSM” to “Disable”.
4. Press F10 to save and exit the BIOS/UEFI.
5. Press Del on reboot to reenter BIOS/UEFI.
 - a. If you are using a SSD-attached via PCI-E, proceed to Step 6.
 - b. If you are using the motherboard’s M.2 slot, proceed to Step 8.
6. If you are using a SSD connected via PCI-E (e.g. an Intel 750) or through an adapter that connects a M.2 SSD to PCI-E, go to “Advanced – PCI-E Configuration” and verify that the device shows on the slot you have it installed.
 - a. Verify the lane count and PCI-E Generation. It should state “x4 Gen3”.
7. Next check the Dashboard on the upper right. The populated slot, lanes used, and PCI-E Generation should all match the information found in the previous step.
 - a. Proceed to Step 9 when done.
8. If you are using a SSD connected to the motherboard’s M.2 slot, re-enter the BIOS/UEFI and go to “Advanced – Onboard Device Configuration,” and set “M.2 Socket3” to “Enable.”
9. Go to the “Boot” Section, set “Boot Mode Select” to UEFI, and set first boot device to “Hard Disk:Windows Boot Manager”.
10. Press F10 to save and exit. Insert/Connect your Operating System install media and reboot.
11. Begin the Windows installation. During the drive selection step, you may need to load additional drivers that are provided by the SSD’s manufacturer, which would be covered in the SSD manual. If these steps are not followed you will likely be unable to install the Operating System to the SSD and make it bootable.

Internal RAID Controller

This section introduces RAID, RAID levels, and the basics of the controller integrated into the PCH. It covers the basics of what RAID does, how RAID works, and why you may or may not want to use RAID.

This section is GREATLY simplified; those who have experience with RAID - especially from the server side - will find this a VERY rudimentary simplification of a process that is quite complex on the back-end. The descriptions of the RAID levels are not based on the back-end logic from the controllers, but rather just a breakdown as a visual reference to understand the basics.

For the sake of explanation in this section, every drive representation here is assumed to be 1 Terabyte (for ease of math); also the reference of P-Drive is “Physical” drive, and L-Drive is referencing “Logical” drives; physical drives are the physical drive you installed, whereas the logical drive is what Windows sees in Disk Management. This is a necessary distinction because RAID is a form of drive virtualization; taking several physical drives and making one logical drive out of them. You can add a minimum of two drives up to a maximum of six drives depending on the configuration (for this controller specifically) and the operating system will see only one logical device.

RAID LEVELS

ALL RAID LEVELS: All RAID levels will lose some capacity through the process of making the array. A small amount of disk space is used, in part, to maintain connections to data across drives, but the quality and size of drives used can impact the amount.

A general rule of thumb for final array capacity is this: Take the number of drives, the array type, and individual drive capacity (for this RAID breakdown all drives are 1TB) and multiply it out as shown below for the specific RAID type you intend to use:

RAID0 = (Number of Drives) * (Drive capacity)

RAID1 = Capacity of one drive, as all data is copied on both drives

RAID5 = (Number of drives – one drive) * (Drive capacity)

RAID10 = (Half the number of drives) * (Drive capacity)

To account for capacity lost both due to maintaining the array as well as what is lost during the formatting process, multiply the product by .85 and you should see a volume close to this number once the array is partitioned. This will vary based on type of drive, its quality, and many other factors; but the number should give you a ballpark estimate on what to expect as a final capacity once formatted.

Please see below for examples of what to expect when you build an array of each type. Each RAID type will use diagrams to illustrate functional and non-functional forms of each RAID; all green items indicate a drive that is functional, and red indicates a drive that has failed.

BEFORE BUILDING AN ARRAY, BE SURE YOUR SSDs OR HDDs ARE THE SAME MAKE, MODEL, AND CAPACITY. MIXING DRIVES CAN CAUSE ANYTHING FROM ODD PERFORMANCE ISSUES, ARRAYS DE-SYNCING, LARGE SCALE DATA CORRUPTION, AND/OR UNRECOVERABLE ARRAY FAILURE. MAKE SURE THAT YOUR INTENDED HDD SUPPORTS RAID; IF NOT, THE DRIVE MAY CAUSE CONSTANT DESYNCRONIZATION ISSUES DUE TO DATA TTL TIMERS NOT SENDING AN “ALL IS WELL” SIGNAL WITHIN THE EXPECTED TIMEFRAME, WHICH WILL SEND THE CONTROLLER THE MESSAGE THAT THE DRIVE HAS FAILED OR DISCONNECTED.

If you are unsure about any of the bolded section above, please contact the drive manufacturer’s customer service to make sure the device is supported for RAID. If the drive is not, the controller cannot work around it and make it function as intended.

RAID0: This type of array is often referred to as “Striping” or a “Striped Array.” RAID0 takes a data set and spreads it equally across two (2) or more drives. The logic behind this array is that reading a single file, for example, will be much faster if the file is spread across two (2) or more drives and read from both drives simultaneously; thus reducing the time each drive spends in the read process by at least 50% rather than being read from one source. The file is then reassembled once the data hits RAM. This is similar, in theory, to how multi-channel memory or SLI works: load balancing for storage.

Because RAID0 is only designed to distribute the data being written across multiple devices to improve performance there is ****NO FAULT TOLERANCE****, meaning if one drive fails, the array fails. It MAY be possible to recover the data but that usually requires a data recovery service, which is not guaranteed and is usually very expensive.

RAID0 is typically only limited by the controller; however, you will get severely diminishing performance returns after 4 drives. If you go above 4 drives, you will also see more capacity loss after creating the array.

The Good-




- it can greatly increase the read/write speed of drives
- The most storage capacity-friendly use of drive space compared to other array types.

The Bad

- RAID0 has no redundancy, which makes it very volatile, if you lose even 1 drive, you lose everything.
- If you use this method, make sure to back up often.

Below is the breakdown of RAID0's function and the breakdown of a drive failure in this type of array. Wherever you see the section labelled "DATA" at the bottom turn **red**, this indicates an array failure.

RAID 0 (2 Drive)



L-DRIVE = $\approx$ 2TB

As RAID0 has no fault tolerance, even a single drive failure compromises the array and renders the array Failed and unrecoverable.

The array depictions below show how this issue scales to larger arrays. In fact, due to the lack of fault tolerance, the potential failure rate actually increases because of the addition of more drives that can physically fail.

RAID 0 (4 Drive)

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-D
DATA-ABCD			

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-D
DATA-ABCD			

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-D
DATA-ABCD			

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-D
DATA-ABCD			

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-D
DATA-ABCD			

L-DRIVE = $\approx$ 4TB

RAID1: This type of array is referred to as “Mirroring” or a “Mirrored Array.”

RAID1 uses two (2) physical drives and writes ALL data to both drives simultaneously providing a 1 to 1 mirror, giving you 100% redundancy live. So as data is being written, if a drive fails you are still up and running. In most cases, when a failure occurs you will experience a stutter in performance and a small but noticeable slow down. Next, you’ll likely see a popup warning fromIRST alerting you that a drive has failed or is disconnected, and your array’s status has been changed to “Degraded”.

RAID1 (at least on these PCH driven controllers) are limited to 2 drives. Also being that this is a mirror, you will use 50% of your capacity in redundancy.

The Good-


- RAID1 allows you to suffer a catastrophic failure of 1 drive with no ill effects to the data being stored.
- Because data is stored on 2 drives at once, read speeds typically increase a little, but not to the speed of RAID5 and, certainly, not to the speed of RAID 0.

The Bad-

- RAID1 is not a storage capacity-friendly array, because the capacity will be limited to 1 drive.

- Due to the capacity available on modern drive solutions, this issue may not be as significant as it once was.
- Write speed will be a bit lower than a single drive.

In the charts below, you can see the effect of fault tolerance when using a mirror array: because all data has a direct 1-to-1 duplicate on the mirrored drive, you can suffer a catastrophic failure of a drive and still retain your data. For a RAID1 array to lose its data, both drives must fail.



RAID5: RAID5 is a stripe with Fault Tolerance, which attempts to bridge the gap between speed and redundancy. This level will always reserve a capacity equivalent to one drive for fault tolerance, regardless of the overall capacity. This means that if you use four 1TB drives to create your RAID5, you will only have the capacity of three 1TB drives; likewise, if you use five 1TB drives to create your array, you will only have the capacity of four 1TB drives. RAID5 requires a minimum of three drives, and the maximum is set by the RAID controller; this level works well when using between four to six drives, but sees diminishing returns beyond six.

Similar to RAID1, or any other current type of array with fault tolerance, a RAID5 array is still usable even while it is experiencing a missing or failed drive resulting in the array functioning in a degraded state. Performance will suffer in a degraded state until the missing drive is replaced and the software rebuild process is completed. During the rebuild process, performance will be severely degraded and can take several hours, depending on the size of the array and percentage of capacity used.

The Good-

- Most space efficient array-type that also has fault tolerance.
- Initial array builds quickly
- Read/write speed is very good; faster than RAID1 but slower than RAID0

The Bad-

- Rebuilding an array can take a long time on an integrated PCH controller, especially when using a very large array with multiple drives or capacity.
 - For this reason, RAID5 is sometimes more beneficial on a dedicated RAID controller.
- If more than one drive fails at once (not likely, but still possible) you will suffer a total loss of data; the array will behave like a RAID0 losing a drive.
- RAID5 is close to outliving its usefulness because it was created, in part, to create large size arrays with fault tolerance. Due to the capacity available on modern drives, other RAID solutions are now better and with fewer downsides.

RAID5 protects data rotating parity (there are several terms coined by different RAID manufacturers over the years for RAID5 and they all mean roughly the same thing), which means taking small portions of data, duplicating them and putting them onto different drives. When a drive that has failed is replaced, its data is recreated by the remaining drives in the array, which will require a very high volume of small data segments copied back over to the replacement drive while also maintaining the array's index of data. Once the rebuild is complete, current data can be read from the new drive and new data can be written to it.

The RAID5 diagram below is a bit more complicated than the others, due to the nature of how data is distributed. In the diagram you will see that there is one more Physical drive than there is data set. This method of data distribution shows that as long as you have three (3) copies of each data set (Data-A, Data-B, and DATA-C), the array will be functional and capable of rebuilding when you add in replacement drive in. And while this is not mathematically correct for HOW the data distribution works, it is a good visualization to understand the basics of how it works.

RAID 5 (4 Drive)

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

L-Drive = DATA-ABC

L-Drive = ≈ 3TB

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

L-Drive = DATA-ABC

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

L-Drive = DATA-ABC

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

L-Drive = DATA-ABC

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

DATA-ABC

P-DRIVE1	P-DRIVE2	P-DRIVE3	P-DRIVE4
DATA-A	DATA-B	DATA-C	DATA-A
DATA-B	DATA-C	DATA-A	DATA-B
DATA-C	DATA-A	DATA-B	DATA-C

L-Drive = DATA-ABC

RAID10: RAID10 is the only level of nested RAID currently supported by the integrated PCH controller. Nested RAID is the process where a controller makes multiple RAID arrays, treating each of them as their own device and applies them to ANOTHER RAID level. This type of virtualization allows you to have an array where each node has its own level of redundancy.

In the instance of RAID10, it is two (2) or more Mirrored arrays (RAID1s) that are striped together (RAID0). In effect, this is RAID0 where every physical drive (now logical drive due to nested RAID functions) in the array is internally redundant via a mirror. This means you can lose a drive and the data set is still present; when the drive is replaced, it is replicated by the data internal to that node.

RAID10 is highly scalable. RAID10 always scales in two (2) drive increments, starting with a minimum of four (4) drives. This motherboard series supports a four (4) drive or a six (6) drive RAID10 array. Lastly, RAID10 is a mirrored array, which means it shares the same 50% drive capacity, as all data has a 1:1 copy.

The Good-

- Excellent Fault Tolerance
- Good rebuild times
 - Vastly superior rebuild times compared to the previous 0+1 methodology of nested RAID.
- Overall performance is good; comparable or slightly faster than RAID5.

The Bad-

- Low space efficiency.
 - With the capacity of modern drives, this will likely be less of a consideration than in previous years.

The fault tolerance works a bit differently in RAID10 than in others because the array has nodes with internal redundancy. In other words, although a four drive array can lose

two drives and remain operational, it greatly depends on WHICH two drives fail. If both drives from the same node fail, then half of the data is gone and the array has failed. However, every node, regardless of the number of total nodes, can suffer one internal failure with no adverse effects.



While the X299 FTW-K controller will support a four or six drive RAID10 array, RAID10 can scale indefinitely provided the controller supports more drives. Every pair of drives adds an additional mirrored node, which increases the theoretical number of failures the array can suffer before a loss of data occurs. However, an array can still fail due to both drives on a node failing.

In the case of a drive failure, the array controller will notify you. When you replace a failed drive in the RAID10, rebuilding the array is mostly seamless. For example, look at the array on the second row to the right. P-Drive1 failed, but P-Drive2 is still working and uses the same data. The array will pull data from P-Drive2 during the rebuild, so the array can be used normally while P-Drive2 copies ALL of its data back to the drive replacing P-Drive1. The rebuild process will only rebuild 1TB worth of data because only one node failed. There will be a performance hit during the rebuild process, which can be further delayed if VERY data intensive applications are used, but overall performance of the array will still be more than fast enough to run effectively during the rebuild. RAID10 rebuilds much more quickly than its predecessor RAID0+1.



RAID0+1: RAID0+1 is a form of nested RAID that was widely used on previous generation boards. Although the X299 series motherboards do not use this type of array, it is listed here to show the improvements made by RAID10, and to clear up a common misperception that RAID0+1 and RAID10 are the same.

A RAID0+1 array is created from two (2) stripe sets that are mirrored together. Similar to RAID10, RAID 0+1 requires a minimum of four drives, and is highly scalable in two drive increments. Again, because RAID0+1 is a mirrored array, it shares the same 50% drive capacity, meaning that four 1TB drives in RAID0+1 will result in a 2TB array.

Where 0+1 differs from 10 is in how the drives are split, and the data distributed. While RAID10 is created using two or more mirror sets striped together, RAID0+1 is two striped sets mirrored together. When scaling with additional drives (in multiples of two), RAID10 adds the drives as another mirrored set to the striped array, whereas RAID0+1 splits the drives between the two stripes to maintain the mirror. To the end-user, the final result appears very similar; however, the significant differences lie in fault tolerance and recovery.

In a RAID0+1, ANY drive failure results in half of the array becoming effectively failed. If one drive fails, that stripe fails, and the mirrored stripe takes over. When the failed drive is replaced, the entire capacity of the mirrored array must be rewritten to the failed array, rather than one drive's worth of capacity (i.e. RAID10). This makes the RAID0+1 array more volatile than RAID10, despite being fault tolerant, and can also increase rebuild times at an exponential margin for large arrays.

Like RAID10, RAID0+1 can afford to lose up to half the number of drives in the array and still be protected; however, this is contingent on the failed units being all from the same stripe set. If one drive fails from both stripe sets at once, the entire array is lost.


The Good-

- Fastest of the standard nested RAID types
- Performance scales with drive count.






The Bad-

- Build times can be substantially longer than RAID10 due to the volume of data being moved, and is typically close to the rebuild times of RAID5.
- Low space efficiency, only 50% of total drive capacity is usable in the array.
- 1 drive failure drops a full stripe set.

RAID 0+1 (4 Drive)



L-DRIVE = ≈ 2TB



Motherboard controllers that support RAID0+1 (such as on older generation EVGA motherboards) will generally support 4 or 6 drive arrays of this type; other controllers can allow this array type to scale indefinitely. Each pair of drives adds to the drive count for the stripes and increases the theoretical volume of failures the array can suffer before a loss of data occurs. However, this type of array can fail due to any two drives on different stripes failing at once; this is RAID0+1's main drawback. Theoretically, at six drives and above, RAID0+1 should have slightly faster read/write speeds compared to RAID10 because the stripes are larger without the overhead of an internal mirror. This is because RAID10 increases both the number of mirror sets as the array scales upwards and the backend calculations needed to maintain the arrays, whereas RAID0+1 only increases stripe size.

As you can see, the difference between RAID0+1 and RAID10 is significant when looking at how data is stored. Although the drive volume scaling and the level of fault tolerance is the same, internalizing the redundancy can make a significant difference overall to the array.

In the examples to the right, you can see that when one drive fails the entire stripe set fails; for a RAID0+1, you would need to rewrite 3TB worth of data back onto the failed node when rebuilding, rather than 1TB for the same drive count on a RAID10.

RAID10 is the current standard on Intel PCH based RAID controllers, largely because the fault tolerance for it is a bit more forgiving and the rebuild speed is overall significantly faster than its RAID0+1 predecessor.



Which types of RAID can I use with my setup?

- 1 Drive – No RAID arrays are supported
- 2 Drives – RAID0 for speed (do regular backups) or RAID1 for data protection.
- 3 Drives – RAID0 for speed (do regular backups) or RAID5 for speed and protection.
- 4 Drives – RAID0 for speed (do regular backups), RAID5 for speed and protection, RAID10 for the best balance of Speed and protection.
- 5 Drives – RAID0 for speed (do regular backups), but it will start to show diminishing returns on performance; and RAID5 for speed and protection.
- 6 Drives - RAID0 for speed (do regular backups), but very marginal performance difference over 4 or 5 drive stripes; RAID5 for speed and protection, or RAID10 for the best balance of speed and protection.

Also, you can run more than one array on your controller, so long as the total is under six (6) drives.

Configuring the Array

Attach all SATA devices you intend to use, and make sure power is attached. Power the system on. Press the “Delete” key repeatedly to enter BIOS. Once into BIOS you will need to enable the RAID function of the board. Once into BIOS, click on (or navigate with your arrow keys) the “Advanced” tab at the top, and then on “SATA Configuration.”



In the “SATA Mode Selection” at the top, the default will be AHCI. Click on the arrow to the right of AHCI or navigate to it with your keyboard and press “Enter” to open the pulldown menu. Select RAID from the list.

RAID mode not only includes the RAID controls, but also shares the same options/functions/commands as AHCI; you may continue using your AHCI devices normally when the SATA Configuration is set to RAID mode.

The SATA Information menu shows a list of all drives currently detected by the controller; make sure the list in your BIOS matches what you have plugged in. If it does not, reseal cables on the device(s) in question. If the issue persists, please check the troubleshooting section on Page 132.



Once RAID is enabled and all devices are detected, press F10 to save and exit. Upon reboot, repeatedly press “Delete” to reenter the BIOS. Once in the BIOS, go back to the “Advanced” tab where you’ll find a new item at the bottom of the list called “Intel® Rapid Storage Technology.”


The Intel® Rapid Storage Technology utility in the BIOS replaces the Intel RAID manager that was previously launched outside of the BIOS in previous generation motherboards. Due to the current bootup process, speed of processors, and UEFI the previous method made accessing the Intel RAID manager nearly impossible. This is the same controller that was previously accessed by CTRL-I; now, it is fully accessible within the BIOS/UEFI.




Once in the RAID controller, you will see a list of all detected drives and a “Create RAID Volume” button. To begin, click on “Create RAID Volume” or navigate to the button and hit “Enter.”




Choose a name for the volume. The controller allows up to 15 characters; you can use numbers and letters, but not special characters.




Next, select your intended array type. This can be done by either clicking on the down arrow and clicking on the RAID level you want, or pressing the enter key and using the down arrow to select the RAID level and pressing Enter again. Please see the top half of Page 64 for a quick reference on different RAID levels and RAID types based on your total number of drives.



Next select the drives you want to use for the array. Select the down arrow and the “X” for each drive you want to include in the array.



Strip size (also called “block size” in other controllers) can be selected manually at 16k, 32k, 64k, or 128k. The controller will determine the default strip size after looking at your drives and array type. Although there are some limited instances where this must be set manually, it is highly recommended to leave this at default.



The controller defaults the capacity to the maximum available space for the RAID. Leaving the capacity at default is recommended because reducing the size is not beneficial, except in limited cases.

To complete the setup process, please select “Create Volume” at the bottom of the page. Depending on the array type, drive type, and volume this can take a few seconds to a few minutes.



Once the array is completed you will see the text shown on the next page, or something similar based on the array type and drives used. At this point, the array is ready to use.



If your array will be your boot drive, the operating system will normally detect the array and see it as a single drive (this is expected), it *MAY* detect it as a RAID array; either way, the OS installation will show the size of the array, not a single drive, and allow you to install the OS to the array without any further steps. However, depending on the OS version, you may need to install RAID drivers for the RAID array to be detected.

If this is meant to be a secondary array, your next step is to partition and format the array within Windows. Please see Page 82 for setup in Windows 10.

Repairing an array within UEFI

This guide will show you how to repair a degraded array from within the UEFI. For testing purposes, a drive was intentionally removed from a RAID5 array and wiped to guarantee that the array rebuild behaved the same as if a new replacement drive was added to a degraded array.

If a drive fails and the array becomes degraded, you will typically see an alert in Windows, during the controller booting (Legacy mode only), or while checking the array status within UEFI. If a drive fails in an array with redundancy (a RAID5 array, in this case), the controller will report the array as “Degraded” (see pic below). This means that the array has a node down, the data is still intact, but your fault tolerance is reduced.




Highlight the degraded array with your mouse or navigate to it with your keyboard and select it to bring you into the array screen (see image on next page), which will show you your array status and any drives that are not currently configured in a RAID array. The Non-RAID Physical Disks list will display any remaining drives on the controller, whether it is a random storage drive, a boot drive, or a replacement drive installed to replace a failed unit. For this example, you will see a degraded array and a “Non-RAID Physical Disk,” which we’ve attached to repair the degraded array.

The drive attached is an exact match to the other drives in our RAID 5; however, you can use a different drive if it is the same size and preferably same type/series, as well. ****ALL DATA ON THAT DRIVE WILL BE REMOVED AND WILL NOT BE RECOVERABLE WHEN ADDING IT INTO THE ARRAY**** The controller will also give you this information, but it cannot be overstated that using a drive with data on it will result in the total loss of all previous data in favor of the data on the array.



Select the degraded RAID volume, then select “Rebuild” on the following menu.



Next, you will see a list of all attached HDD/SSDs that can be used to rebuild the array. Select the disk, then click on it or press enter.



Once the process has started you will see the status change to “Rebuilding.”




****Important Notice**:** The controller will not begin the rebuilding process until you have booted back into Windows; this queues the rebuild but does not start the process.

The rebuilding duration will vary. Fast drives in uncomplicated arrays with small amounts of data can rebuild very quickly. Large platter drive-based arrays, especially in RAID5, can take several hours to rebuild. The duration of the rebuild process will further increase based on the usage of both the array and the CPU during the process.

The rebuild process can be monitored from the “Intel Rapid Storage Technology” utility in Windows, which is effectively the software front-end for the Intel SATA controller. The rebuild status can be viewed in the “Manage” tab, located on the top of the IRST Window. Please see the image at the top of Page 81 to see an example of where to find the rebuilding % in the IRST.


IRST (Intel Rapid Storage Technology)

The IRT is the software front-end for the Intel SATA controller. It is recommended to install the IRT drivers after installing the Intel Chipset Drivers – the main motherboard drivers. This guide walks you through not only the building, but also the repairing of an array from within Windows, since it may be more convenient to build new arrays within the OS rather than from within the UEFI. All forms of RAID that can be built in the UEFI are available in IRT; the IRT may be more convenient for some people due to a more detailed UI than the UEFI.



This guide will illustrate how to create a RAID1 array using the IRT. Using the screen above, we know that three (3) SSDs are currently attached to the controller and the other three (3) INTEL SATA ports are empty (this can vary board to board based on how many Intel SATA ports are present). If you plan to create another type of array, please ensure that you have the correct number of drives attached and shown on the status screen in IRT for your intended array before you start.

Referring to the above image, again, locate and click the hyperlink halfway down the window that says “Create a Custom Volume.” An array creation window will open:



SATA will be selected by default. PCI-E primarily refers to PCI-E / M.2 based NVMe drives; the same basic steps do apply to both, however. Select SATA, and “Real-time protection (RAID1).” Then, click Next at the bottom of the window.

RAID1 can only support two (2) drives. Left-click the open boxes from the list below to select your two preferred drives. You may also name the array anything you like. Although you have the option to define volume size, it is recommended to leave the array at 100%, which is the default.

Create Volume...

1. Select
2. Configure
3. Confirm

Configure Volume

Name:

Select the array disks (minimum selection required):

- ☒ SATA SSD on Controller 0, Port 0 (112 GB)
- ☒ SATA SSD on Controller 0, Port 1 (112 GB)
- ☐ SATA SSD on Controller 0, Port 3 (112 GB)

Volume Size Advanced

Volume Size 114,473 MB

Array allocation:

Proposed Configuration

Array_0000

Volume_0000
Type: RAID 1
112 GB


Back **Next** Cancel

[More help on this page](#)


In the Advanced tab, you can select the option to “Initialize Volume,” which will occur after the array is created. If the array is not initialized now, it can be initialized later in “Disk Management.” See Page 82 for Disk Management instructions.

When done, click “Next” at the bottom.


The screenshot shows the 'Create Volume...' wizard in the 'Configure Volume' step. On the left, a sidebar lists three steps: '1. Select', '2. Configure' (which is highlighted), and '3. Confirm'. The main area is titled 'Configure Volume' and contains a 'Name' field with 'Volume_0000'. Below this, a section 'Select the array disks (minimum selection required):' lists three options: 'SATA SSD on Controller 0, Port 0 (112 GB)' (checked), 'SATA SSD on Controller 0, Port 1 (112 GB)' (checked), and 'SATA SSD on Controller 0, Port 3 (112 GB)' (unchecked). A 'Volume Size' section has two tabs: 'Standard' and 'Advanced' (which is selected and highlighted with a red box). Under the 'Advanced' tab, there are two checkboxes: 'Enable volume write-back cache' (unchecked) and 'Initialize volume' (checked, also highlighted with a red box). To the right, a 'Proposed Configuration' window shows 'Array_0000' with two 112 GB disks and a volume 'Volume_0000' of type 'RAID 1' with a size of '112 GB'. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red box. A link 'More help on this page' is visible in the bottom right corner.



Review the summary provided on the confirmation screen. If you are unsure about any selections made, click the “Back” key and make your corrections. When ready, click “Create Volume” at the bottom. This typically takes between a few seconds to a couple minutes depending on the size and complexity of the volume. Once finished, you will see the message below. The array is now ready to be partitioned, formatted, and used.




Once you click the OK button on the RAID creation window you will be brought back to the main window, “Status” tab. If the option to initialize was selected, the initialization status will be shown below, circled in red.



If you select the “Manage” tab, you can see a bit more in-depth information on the array and additional options to manage or change the array.

Before the drive can be used, however, the drive must be partitioned, formatted, and assigned a drive letter. Please see Page 82 for instructions.




Repairing an array within IRST

This section of the guide will illustrate how to repair a degraded array from within the IRST. For purposes of this guide, we are repairing a degraded RAID 1 array using a third drive plugged into the controller, but not currently in use.


Below, you can see a degraded array, and one of its drives reported as missing/failed. The data on the other drive is still intact, but the fault tolerance is offline due to the missing/failed drive. The IRST also shows several warnings, the "!" for "Status," a "!" for the portion of the array missing (showing the degradation of the logical drive), and the "!" for the physical drive, which also states 0GB – a further indicator of a faulty drive. You can repair the array from here or from the Manage tab at the top of IRST.




The “Manage” tab shows the array specifically, and not just the controller as a whole.




Next to “Status: Degraded,” left-click the hyperlink labeled “Rebuild to another disk.” This will bring a pop-up window over the IRST showing a list of attached drives that can be used for the repair (see pic on next page):




Select the drive you wish to use for the repair and click the “Rebuild” button.






The rebuild process will begin. As with any RAID array with Fault Tolerance, the rebuilding time depends on several factors, such as array size, array type, CPU, etc. You will then see the Rebuild % status in the Manage tab. Once repairs are complete, the array will update to “Status: Normal.”



Partitioning and Formatting a drive

Once you have created your array, either from UEFI or from IRST, you will not initially see your array in “This PC.” This is expected, because even though you have created the array, you have not yet prepared the array to be used.




To begin, you’ll need to go into Disk Management.

Windows 8, 8.1, and 10: Right-click on the Windows Start button and select “Disk Management.”

Alternatively, press the Windows Key + X on your keyboard and select “Disk Management.”


Windows 7: Left-click the Start menu, right-click on “Computer” and select “Manage” from the drop-down. The Computer Management window will pop-up, and you’ll select Disk Management from the menu on the left, under Storage.



After “Disk Management” loads, you’ll see a pop-up to Initialize Disk if you’ve added a new drive or created a new array.

Generally, it’s recommended to select “GPT,” unless you need backwards compatibility with an old OS or PC. When you’ve made your choice, click “OK.”


Note: If you previously initialized your array through IRST, you will skip this step and move on to the New Simple Volume Wizard in the next step.



Before you can assign a drive letter to a drive or array, the initialized disk must be partitioned. If you are following this guide and just initialized your drive or array, the New Simple Volume Wizard will automatically pop-up.

If your drive or array is initialized, but not partitioned, the disk will appear in Disk Management as “Unallocated,” as shown in the image below. Right-click on the box containing “Unallocated” and select “New Simple Volume” (this text may vary slightly based on operating system). The New Simple Volume Wizard will pop-up.

When the window below opens, click “Next.”




Leave the size at default to create a partition using the entire volume of disk space, then click “Next.” Select the drive letter you want to represent this drive, then click “Next.” Note: The drive letter does NOT have to be a consecutive letter with previous drive(s). Also, the wizard will not allow you to accidentally select the letter of a drive in use.

The image shows two side-by-side screenshots of the 'New Simple Volume Wizard' window. The left screenshot is titled 'Specify Volume Size' and shows fields for 'Maximum disk space in MB' (343283), 'Minimum disk space in MB' (8), and 'Simple volume size in MB' (343283). The right screenshot is titled 'Assign Drive Letter or Path' and shows the 'Assign the following drive letter' option selected, with 'D' chosen from a dropdown menu. Both screenshots have red boxes highlighting the 'Next >' button at the bottom.


The next step is to format the partition. Select your File System; NTFS is default, and generally recommended for most large drives. Check the box for “Perform a quick Format.” You may rename the volume, or leave it at default. Click “Next” when ready.

The image shows a screenshot of the 'Format Partition' step in the 'New Simple Volume Wizard'. It prompts the user to choose whether to format the volume and, if so, what settings to use. The 'Format this volume with the following settings' option is selected. The settings shown are: File system: NTFS, Allocation unit size: Default, Volume label: New Volume, and 'Perform a quick format' checked. The 'Next >' button at the bottom is highlighted with a red box.

After the quick format is completed, you will see the last Window of the wizard, a summary of the process, then click “Finish.” The drive is now usable.



To confirm, go back to File Explorer in Windows. Click on “This PC” and check the drives section. You should have a new empty drive there, with the letter you designated.




At this point the process of building an array and making it usable is completed.

Fan Header DC and PWM setup

This motherboard supports both 4-pin PWM fans and 3-pin voltage-controlled fans. All of the fans can be set to a static voltage manually. The upper two (2) fan headers (see Pages 12 and 13, component number 3) are PWM-controlled, whereas the remaining fans (see Pages 12 and 13, component number 4) are controlled by either DC (Direct Current) or PWM.



To configure the fans in BIOS/UEFI, first power on / restart the PC. During the POST sequence, press [Delete] repeatedly to get into the BIOS. Once in the BIOS, use the arrow keys or your mouse to navigate, whichever is easier, and make your way to the “Advanced” – “H/W Monitor Configuration” menu.



Once into the H/W Monitor section, you can see the temperature monitors across the top. Below the monitors are the two (2) PWM fans: CPU1_FAN, CPU2_FAN, and CHA_FAN. The PWM fans are set to “Smart,” which means the PWM controller is using a Smart curve for fan controls. Each PWM fan can be set to a separate fan curve. To set a Smart curve, please see the guide starting at the top of the next page.

The +/- buttons next to Smart allow you to set your fans at a static speed percentage, instead of a curve. If you set the speed too low, however, the fan may stall; the stall speed will vary from fan to fan.

Below the fans, you will see the four (4) DC fans: SYS1_FAN, SYS2_FAN, PWR_FAN, and AUX_FAN. These fans can only be set to a static speed percentage. Again, if you set the speed too low, however, the fan may stall; the stall speed will vary from fan to fan.


The DC/PWM fan headers can be set to either DC Mode or PWM Mode, which can be selected by clicking on the pulldown menu. DC Mode will power the header using the static percentage set and the fan’s maximum speed. PWM Mode will power the header and communicate via PWM with a supported PWM fan. PWM mode is necessary if you are running a PWM fan powered by a different connector, such as a molex, SATA, or 3-pin header not on the motherboard. If you don’t select PWM mode in this scenario, the fan may run at 100%, or behave erratically, regardless of the static percentage you set above. If you have a PWM fan powered directly by the 4-pin header, then selecting PWM will control the fan via PWM signal NOT via reducing voltage, while it will run at a static speed based on the percentage selected, this also should reduce the low RPM stall point of the fan as it is receiving full voltage. A DC fan will run at its static percentage regardless of whether you set DC Mode or PWM Mode.

To set a Smart curve, select the “Smart Fan Settings” and enter the menu.



First, choose the temperature monitor the PWM controller will use to monitor for its temp information. It's recommended to tie the fan control to the CPU, which is predominantly the most important temperature in the system. The exception is when you are pushing memory overlocks HARD with 64GB+ of RAM installed, which may cause PWM temps to be a concern; in that case, tie the fan control to the PWM temp.

Once you've set the Fan Control reference point, you can set the Default fan speed. The fan speed will increase once the temperature reference point has heated up enough to hit the Level 1 Temperature. Once it hits the Level 1 Temperature the fan controls will override the Default speed setting based on the temperature at the time.



There are four tiers of temp control, with Level 4 fan speeds generally recommended to be set at Max. The example above uses a fairly aggressive fan curve, but you can set this as high or low as you wish. Make sure, however, to stay below 5-10C of the Max safe load temp for your specific processor, which can be found at Ark.Intel.com. All Smart Fan Settings have the same controls and can be setup the same way.

When monitoring temperatures vs. fan speed, you may notice a variance in ramp up/down temps; this is due to a function EVGA hardcodes into the BIOS called Hysteresis. Hysteresis builds in a buffer to control fan speed behavior. This feature prevents a constant ramp up/down from happening when your system sits exactly at the temp you set for SMART fan controls. For example, if your setting is at 30C and you hit 31C, the fan will ramp up and cool down to 29C, letting the fan slowdown, which lets the system heat back up to 31C again, repeating indefinitely and causes the system to sound like it's breathing. Hysteresis adds a 4C +/- buffer for CPU and PWM, and 2C +/- for SYS. This means that once a fan based on the CPU or PWM temp brings you down to 30C, it will not ramp back up until 34C, so you have a buffer. If you base the fan control on the SYS temp, you will only see 2C variance.

Setting Up SLI and PhysX

SLI and PhysX are NVidia developed technologies supported by EVGA motherboards. SLI is a means of dividing the graphical load between multiple video cards, similar in theory to how dual- or quad-channel memory works, or even how RAID0 works. However, the software (e.g. game) needs to support SLI to gain from this technology.

PhysX is a NVidia technology that can be assigned to a graphics card or CPU for programs supporting NVidia PhysX. As long as the PhysX driver is installed, which occurs during a normal NVidia GeForce driver installation, the system will allocate resources to support PhysX based on CPU or GPU load. You may, however, go into the NVidia Control Panel and dedicate the CPU or a GPU to this function manually. For a complete list of NVidia PhysX supported titles, please see this link:
<http://www.geforce.com/hardware/technology/physx/games>.

SLI requires two or more video cards that support SLI; have the same GPU family (GTX 1080 Ti, GTX 480, etc.); have the same memory type (GDDR5, GDDR5X, etc.), volume (3GB, 8GB, etc.), and datapath width (128bit, 256bit, 384bit, etc.). SLI also requires a motherboard that supports SLI and a SLI bridge for the gold fingers of each card in SLI. Providing the above conditions are met, compatible graphics cards with different GPU and/or memory clock speeds will not prevent SLI from being enabled.

Note: For questions and concerns regarding NVidia Surround™ and full setup instructions, please see <http://www.geforce.com/hardware/technology/surround/system-requirements>. As the possible combinations of setups between SLI, monitors, and cabling are far too vast to label in this manual, also see <http://www.geforce.com/whats-new/guides/how-to-correctly-configure-geforce-gtx-680-surround> for a basic setup walkthrough.

Installation:

1. Physically install your graphics cards, then install a SLI bridge; examples include a Flexible bridge (included with this motherboard), an EVGA Pro Bridge, or an EVGA HB Bridge. Current NVidia graphics drivers support 400 Series Fermi cards up through GTX 1080 and TitanX Pascal cards. Driver support is determined by NVidia; please check GeForce.com to confirm compatibility with drivers and SLI.



Once the graphics cards are physically installed, connect a SLI Bridge; your cards should look similar to either picture below.





2. After the cards are installed, have power connected, and the SLI bridge attached, boot into Windows. The graphics driver will normally identify the cards and automatically configure the driver. If not, then you may need to reinstall the driver. Once finished, you will receive a popup in the lower right corner stating that you have a SLI capable setup and it needs to be configured. If you did not see the message, then first verify that both cards are detected and functioning without system errors from Windows Device manager.

Right-click on the Start menu and select Device Manager. For Windows 7, left-click the Start menu, then right-click on Computer, select System, and you'll see Device Manager on the upper-left side of the popup window.

Under "Display Adapters" you should see the type and number of video cards you have installed.



3. Once you have verified there are no detection/driver installation issues with the cards you can enable SLI. Right-click on the desktop and select “NVIDIA Control Panel” (“NCP”). Next, select “Configure SLI, Surround, PhysX” under the “3D Settings” menu. The default is “Disable SLI.” To enable SLI, click “Maximize 3D Performance,” circled in red below, and click “Apply” at the bottom.




Before you can enable SLI, the NCP may ask you to close certain programs and processes; you cannot continue further until this is completed.

At this point the display may go black a few times, or appear to change resolution and back again quickly; this is normal. Also, if you are running a G-Sync monitor, it will turn off and on during this process, which may take approximately 5 and 20 seconds.

4. Once finished, your SLI configuration should look similar to the image below. The key aspects that show all is well are circled.

5. OPTIONAL. With multiple cards, you can assign one to PhysX, but we do not advise choosing that OVER SLI, because SLI will provide substantially larger performance gains. In general, leaving PhysX on “Auto-Select (recommended)” is recommended. When left on Auto, your PC will adjust the PhysX duties based on overall system load, which, in many cases, will be the CPU. Otherwise, manually setting the PhysX processor is best when a spare graphics card is in the system.



While SLI is a direct performance enhancement PhysX is used for visual enhancement, allowing the physics of various items in the environment to be calculated live. A dedicated PhysX card keeps the load localized to a device that has no other current function. Although frame rates may only increase slightly, a dedicated PhysX card keeps the frame rate from spiking in either direction when there are very intense and abrupt uses of PhysX. For certain titles that heavily utilize PhysX, a dedicated PhysX card can be beneficial to help stabilize the frame rate.

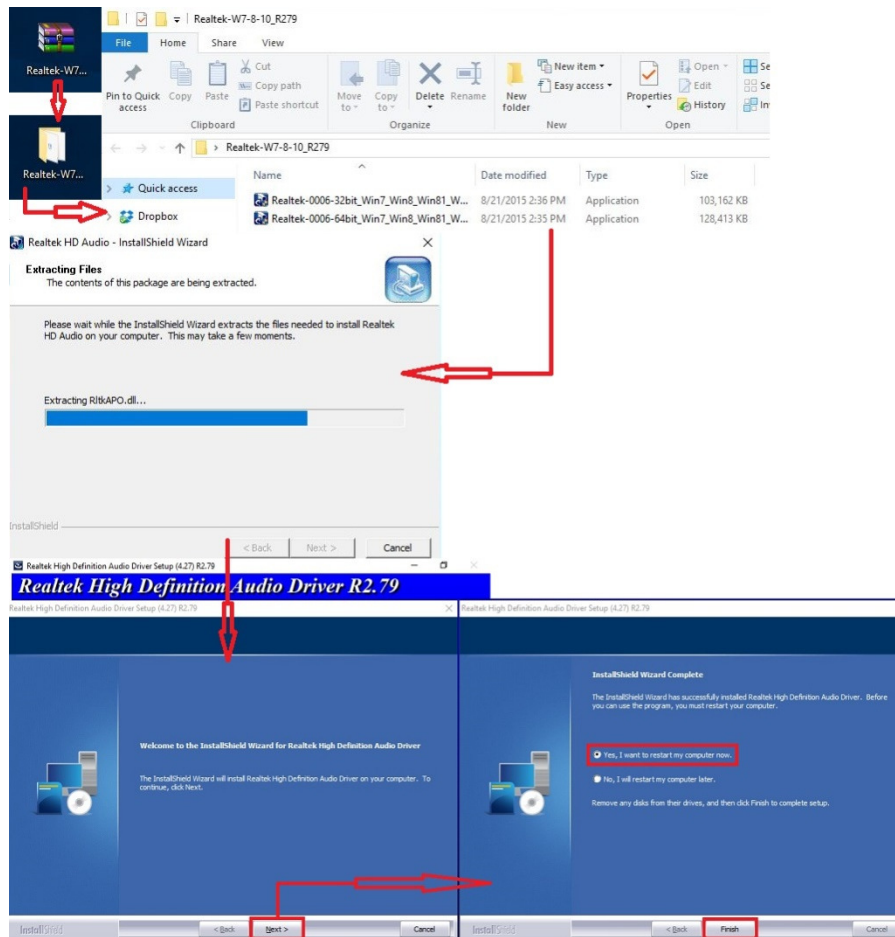
If you do not have a spare card to dedicate to PhysX then it is advised to leave the PhysX Setting to "Auto-Select (recommended)." The general rule of thumb for needing dedicated PhysX is this: Does your game support NVidia PhysX? (Yes or No). If “No,” then a dedicated PhysX card serves no purpose. If “Yes,” then the next step is to see if your GPU has a high usage rate while playing normally. Use a program like EVGA Precision XOC to monitor the GPU usage of all current video cards. If the GPU is consistently over 75% usage, the GPU usage occasionally maxes out and the frame rate drops in moments of intense action, then dedicating a card may be beneficial. However, if this does not occur in your setup, then leaving your system to handle the PhysX load without a dedicated card should be fine, since the default setting will balance the PhysX processing between the GPU and CPU based on load at any given moment.

If you need further help setting up SLI/PhysX, please contact Customer Support.


Realtek HD Audio Manager

The X299 FTW-K and MICRO both use a 7.1 Realtek ALC1150 audio controller. This section will cover installation of the controller (in Windows 10) and the basic configuration options that are available in the software.

To begin, use your installation DVD or download the current Realtek drivers from www.evga.com/support/download. Once the zip file is downloaded, right-click the file to extract it, and run the correct 32-bit or 64-bit executable for your operating system. Once launched, select “Next” on the extraction window, “Next” on the installation window, then select “Yes, I want to restart my computer now.” Click “Finish.”




Once you have rebooted and re-entered Windows, there should be a new shortcut in the system tray with a red-orange speaker icon. When you mouse over it, the icon reads “Realtek HD Audio Manager.” When the menu opens, left-click on it, and left-click the similar red-orange speaker icon labeled “Sound Manager” to launch the Realtek software front-end.



The below images are all composites, showing the popup texts for multiple buttons/menus on screen at once to reduce the picture clutter; however, you will not see the same popups simultaneously in the Realtek HD Audio Manager.


Once the Sound Manager is open, you will be able to configure several aspects of the sound output. First on the main screen is a side-panel showing the devices physically attached to the Realtek controller. See Page 14 for output jack legend.




On the right margin, you can see the different ports labeled in the picture above. The top portion covers the five (5) analog ports on the I/O Hub. The two (2) labeled “Speaker Out” ports and “Line In” port all use a diffused color to show that a port is not currently connected; conversely, “Front Speaker Out (Speakers)” and “Mic In (Microphone)” use a solid color to show that the ports are currently connected. For locations and a picture of the ports please see Page 12 component number 31, and Page 14 for the I/O Hub component number 9.

The center section of the side panel is for Front Panel Audio. Front Panel Audio ONLY supports 2.0/2.1 speakers, or headphones and a standard Microphone. These are only usable if you are using a front panel connector, such as a 5 ¼ bay device with audio connectors or the front panel audio pre-attached to your chassis. The Front panel header can be found in the component legend on Page 12/13 component number 24, additional port information on Page 20, and a header pinout on Page 45.

The bottom section of the side panel is “DIGITAL.” This represents connections via optical S/PDIF, typically going to a receiver or TV. For locations and a picture of the port please see Page 12 component number 31, and Page 14 for the I/O Hub component number 8. The optical out contains all audio data the controller can provide; however, only 2.0/2.1 is pre-encoded from the controller. Audio with a higher channel count through optical S/PDIF must be connected to a receiver that supports live encoding, or an equivalent technology.




Also on the side panel is the “Device Advanced Settings” button. This will bring up a popup menu with options for playback and recording devices. Both options in Playback Device and Recording Device are mutually exclusive within their respective sections. The top options are default for both. Once you have made your selection, click “OK” at the bottom.




The icon of the file folder in the upper right is “Connector Settings.” This is the window for setting up port detection. The default setting in this Window allows the Manager to create a popup notification to let you know when you’ve plugged in a new device into a Realtek port(s) while Windows is running. The popup box will let you confirm which device you’ve plugged in (e.g. after plugging in a Microphone cable, the popup box may incorrectly select “Line In,” which you should then change to Microphone). If you’d prefer to disable this notification, then deselect the “Enable auto popup dialog” option.

The “Disable front panel jack detection” option is mostly used for troubleshooting in cases where the Realtek Manager refuses to acknowledge speakers plugged into the rear ports, causing no sound to come out from your speakers. Should you run into that problem, checking the box for this option will often fix situations where audio will not play. If you plan on using your front panel audio connectors (e.g. the speaker/mic ports on a chassis), then do not disable the front panel jack detection. The disable box *ONLY* applies to the front panel, whereas the “Enable auto popup dialog...” covers ALL Realtek interfaces.

The “i” located at the bottom right of the Manager is an information button. Clicking on this will provide DirectX information, the CODEC and Language selection options.




The main page defaults to the first audio device it detects; device response time to the controller can vary, so the order listed here may not be the order ALL components are detected. The Manager will generally only show devices detected by the motherboard. For purposes of this guide, speakers, a Microphone, and a TOSLINK digital audio cable are connected. As a result, the Manager detects the Digital Output first. You should also see a “Digital Output (Optical)” at the top. This shares all the same options as “Digital Output,” the difference is strictly connection type.



In the “Default Format” section, there are a couple basic selections of “CD Quality” and “DVD Quality,” which allow you to select optimized presets based on standards of how the media is recorded and written to a disk for those formats. Use the pulldown menu if you prefer to set the bit-rate and frequency yourself. All audio devices shown will have a Default Format section. Adjusting the frequency too low will heavily reduce audio quality, moving it too high can impact overall performance and cause other anomalies.

To the right of the pulldown is the Restore Defaults button, shown by an eraser icon, which will reset any changes made to the default of the controller.



If you switch over to the Sound Effects tab, you'll notice the "Main Volume" controls near the top. The slider directly under Main Volume adjusts the Balance of your speakers. It is generally recommended to leave the slider centered, since most audio is meant to run with the audio output centered; this setting is mostly to compensate for either off-center speaker configurations or individuals suffering from partial hearing loss.

To the right of Balance is "Volume;" this is the volume control for the Realtek software but also adjusts Windows main volume; this will also work similarly to a pre-amp if your speakers have a separate volume control that controls the volume after leaving the PC but before reaching the speakers. For this reason, many modern speaker systems use sequential volume controls working additively off each other.

To the right is the Mute button. This is tied-in with the Windows volume; clicking this button will Mute all audio output in Windows. Click again to restore audio.

The "Set Default Device" will allow you to set the Default Device or Communication Device without entering the Windows Playback menu; for non-default devices, this box will drop down similar to the image above. For default devices, the box will be grayed-out if the device selected above (i.e. Speakers, Digital Output, etc.) is the default device.



Below the Main Volume section is the “Sound Effects” and “Default Format” section. “Default Format” was previously covered on Page 100. Under “Sound Effects,” you’ll find Environment, Loudness, and Equalizer options.

Environment is defaulted to “None,” which means the sound will output how it was originally recorded.

This section allows the Realtek controller to simulate different acoustic environments before delivering it to the speakers so the sound will emulate an arena, an auditorium, or even a bathroom.

If you want a preset other than what is available on the main screen, click the pulldown box next to “Environment,” which defaults to <None> and contains a list of 24 total options.

You can revert the sound effects back to a default state anytime by clicking the “Reset” button.





To the right of the Reset button is the “Loudness Equalization” setting. This is effectively volume normalization, which is a setting designed to prevent drastic volume changes in music, movies, games, etc. from one audio source to the next, much like many TV’s utilize to prevent a sudden increase in volume when switching from one channel to the next. This is a good setting if you prefer to avoid volume going higher than you want; however, you may lose some audio quality, depending on the audio source. This is unchecked by default. There is no additional configuration; the setting is either enabled or disabled.

The last section in Sound Effects is the Equalizer. Much like the environment section, there are visual references for presets, a pulldown with 14 options, and a more traditional equalizer setup with a 10-band EQ as part of the software suite.

You can select any of the presets that will provide you with basic optimizations for your music preference. To see the entire list, click on the pulldown menu that reads <None> and a pulldown with all presets becomes available. Including "<None>," there are 14 in total.


If you prefer to set your equalizer values manually, click on the equalizer icon to bring up the 10-band equalizer. This will bypass all presets and enable manual control.



From here, you can adjust the tone of the music by raising or lowering the slider on each frequency band.

You can make multiple presets of your own, which will appear in the pulldown after you have saved them. The presets each have their own benefits, but cannot take into account your audio equipment, room acoustics, or subjective audio preferences. Creating your own EQ profile allows you to further customize your audio to account for additional musical options, such as Thrash Metal, Country music, Opera, etc. to match your specific tastes.

If you would prefer to not manage the EQ in full manual mode, you can click the guitar icon with the red box around it to bring you back to the equalizer presets menu.




The last section in Sound Effects is voice cancellation, also called Karaoke mode (this should not be confused with noise cancellation options in the Microphone menu). By left-clicking the icon, a red X will pop up, as shown in the image to the right. When the red X is present, the Realtek controller uses an algorithm to suppress vocal sounds from the audio source, while keeping instrumental sounds, much like a karaoke machine.



Underneath Karaoke is a menu for changing the key of the music playing. Raising the value (up to +4) will raise the pitch of the music and voice, while lowering the value (down to -4) will lower the pitch of the music and voice. This setting is designed to add a fun effect to your audio; it generally should not be used as a means to “fix” audio.


The next device section is “Speakers,” which is where you can configure all of the output options while using the analog connectors on the back.

This section will focus on the “Speaker Configuration” and the “Room Correction” tabs. The other tabs, “Sound Effects” and “Default Format,” are identical in each section, and were previously covered on Pages 101 and 100, respectively.





The first option in “Speaker Configuration” is the pulldown menu for Stereo, Quadraphonic (4.0), 5.1 Speakers, and 7.1 Speakers; there is no support for 4.1 audio.

The “Audio Test” button will be present for all speaker configurations, and serves the same function in all speaker iterations.



Pressing the “Auto Test” button will play the same sound file out of each speaker that should be present, based on the configuration option you selected. As the sound test plays, an animation will highlight the speaker the sound should be coming from, this will help to ensure that the speakers are connected correctly and oriented correctly, as well.



In the lower left, there is a pair of checkboxes labeled under “Full-Range Speakers.” Full-range speakers are designed using drivers that can independently output audio across a wide frequency response of (20Hz-20,000Hz +/- 3dB) at the same volume level. When this setting is enabled, Windows will disable its software crossover settings and send the full frequency range to the speakers, whereupon the speakers will engage its own hardware crossover filters to appropriately direct audio to the individual drivers.

If your speakers are not full-range, they will not have necessary hardware components to play low frequencies, nor are they designed for that purpose. When the full-range option is unchecked, the Audio Manager will instead send the full signal out using crossover settings, available on 5.1/7.1 by selecting the Bass Management option. This sort of signal will send out the full audio frequency such that the majority of the frequency will go to the speakers, but direct the low frequencies to a subwoofer. This way, you’ll still get the same sound, but you must select the correct options.

To determine if your speakers are full-range or not, you must look at the minimum frequency response of your speakers. The concern is not the high-end of the frequency response, but rather the low-end. As noted above, if your speakers are not rated for a minimum of 20Hz, then they are not full-range speakers. Here are some considerations:


- Assume your speakers are not full-range. Modern speaker design uses a subwoofer to handle the relatively small percentage of content you hear in an audio signal. Low frequency content, more commonly known as sub-bass, is greatly amplified by the subwoofer, which explains why the smallest percentage of your audio content can often be the loudest. The subwoofer allows the speakers to concentrate on the remaining 99%+ of the audible frequency bandwidth.


- Generally, if your speakers came with a subwoofer, they are not full-range speakers. As noted above, subwoofers are designed, in part, to allow speakers to stay clear of the low frequencies. When examining the specs of the speakers, make sure to look at the frequency response specifically of the speakers, and not just the specs of the speaker system. Sometimes, the most easily available specs may include the subwoofer, which can go as low as 18Hz +/-2 dB; the minimum frequency of the speakers, however, may be much higher.
- If you have a 2.0 stereo setup, look for the frequency response specs. If they do not have at least a 20Hz to 20kHz range, then they are not full-range speakers. You may try full-range, but you will lose any audio data below the minimum frequency of your speakers or may hear distortion at lower frequencies.
- Headphones tend to be full-range speakers, as they are designed to operate without a subwoofer. However, the quality of your headphones will determine whether you should configure them as full-range or not.
- If you use a receiver or your speaker system has its own controls for the type of speakers, do not use the full-range setting in the Creative panel without consulting the instructions of your receiver or speaker system, since it's possible that your receiver or speakers prefers full-range, so that the receiver or speaker system will utilize its own crossover settings once receiving the full audio range.
- Lastly, audio equipment quality will have the greatest effect on whether you should set full-range, or not. Very few speakers are properly designed for full-range audio, and if they are, then they will likely be very expensive.

Two last considerations: First, do not set Headphones to full-range if using the Headphone Surround option (may be listed as Virtual Surround), as Surround will not work properly in full-range mode. Second, do not enable full-range mode for 5.1 audio or the bass redirection will not work properly. The only exceptions to 5.1 and full-range is that you either 1) have a full or partial set of full-range speakers in your setup, or 2) you are using a receiver or speaker system that prefers to handle the low frequency crossover at the receiver or speaker system, rather than using the Audio Manager to configure the crossover settings for low frequencies.

When Full-range is enabled, the speaker icons are enlarged to represent the change.

Finally, Virtual Sound is an option created to send a stereo signal with positioning data to a Matrix decoder to simulate surround over a stereo signal.





Quadrasonic and above will have similar options for Full Range: the checkboxes for “Front Left and Right” and “Surround Speakers,” which controls the satellites.

Under Quadrasonic, there is an option labeled “Speaker Fill.” This feature uses an algorithm to manipulate a stereo audio source to create a virtual surround effect using additional speakers.

5.1 and 7.1 add more options. “Speaker Fill” carries over from Quadrasonic, but also includes “Swap Center / Subwoofer Output” and “Enable Base Management.”

Speaker Fill is still a virtual surround sound emulation feature for stereo audio sources, but not a true surround.

Swap Center / Subwoofer output is an option because the Realtek controller shares the subwoofer port with the front (stereo) speakers on the rear panel audio. As the tooltip notes, if low frequencies are not playing on the subwoofer, you must enable this setting. Also, most surround speaker sets use both the front pair and subwoofer on the same channel and voice speaker on a separate channel; in some instances, this is reversed such that the voice channel is joined with the front pair on one channel and the subwoofer has a dedicated connector, which can lead to voice channels and subwoofer receiving the wrong signal. The "Swap Center / Subwoofer Output" can resolve this type of issue. Please first confirm the subwoofer is installed correctly, however, by using the Component Legend on Page 14, and by consulting your speaker set’s manual.

Enable Base Management (on the previous page) is a setting that attempts to keep audio data that may be lost or distorted near the frequency crossover point for the subwoofer and satellite speakers. On a system that does not use full-range speakers, this is a crucial setting for ensuring that your speaker system correctly plays low frequencies. The default frequency is set to 80Hz, but your system may need adjustment.



To properly set the crossover frequency, look at the specs for both your speakers and subwoofer. Some pre-assembled speaker systems may tell you what to set the crossover frequency to in your instruction manual. If not, then look for the minimum frequency response of your speakers (i.e. 75Hz, 80Hz, 175Hz, etc.), and the maximum frequency response of your subwoofer (i.e. 60Hz, 80Hz, 120Hz). Once you have that information, you can set the crossover frequency.

Ideally, you'll want to make sure that the crossover settings provide enough room to set the crossover frequency a little higher than your speakers' minimum frequency response, which will reduce distortion and provide a smoother transition to your subwoofer. A good rule of thumb is to set the crossover frequency about 20Hz higher than your minimum speaker frequency response. However, you may run into three different scenarios:

First, the frequency response of your speakers (e.g. 60Hz) and subwoofer (e.g. 100Hz) nicely overlap. This is ideal. Using the general rule above, you would set the crossover frequency to 80Hz and call it a day. You should still test the audio to ensure it is free of distortions, but this provides enough room to pass the low frequencies to the subwoofer and keep the higher frequencies to the speakers, without creating a conflict.

Second, the maximum frequency response of your subwoofers is exactly the same as the minimum frequency response of your speakers. In this case, you must set the crossover frequency exactly to where the subwoofer and speakers meet.

Third, you may have a situation where your speakers' minimum frequency response (e.g. 100Hz) is higher than your subwoofer's maximum frequency response (e.g. 80Hz). This should only happen if you purchased the subwoofer separately from the speakers. This setup cannot work properly, and you must replace either the speakers or subwoofer. Although you can adjust the setting to either the subwoofer's maximum frequency response or your speakers' minimum frequency response, you'll lose the audio data between those frequencies.

The Realtek HD Audio Manager, unfortunately, does not contain the ability to adjust the crossover frequency settings. Instead, you must adjust the crossover frequency via a registry edit in Windows. Due to the advanced nature of editing the Windows registry, we do not provide instructions in this manual. However, if you would like to adjust the crossover frequency settings, please see our FAQ on the EVGA website to walk you through how to adjust this setting: <http://www.evga.com/support/faq/?f=59663>.


The 5.1 and 7.1 configurations include the section "Optional Speakers." This allows you to disable satellite speakers in pairs to bring any 5.1 or 7.1 system down to stereo. This is mainly used as a diagnostic tool for testing sound issues.



In the image above, speakers become transparent in the image when they are unchecked; this is to give you a visual reference of which speakers are enabled and disabled when testing for sound.


The last section under “Speakers” is “Room Correction.” This section allows you to virtually position the speakers using the software to ensure that the surround effects are positioned correctly based on physical distance.

This is **ONLY** supported in 5.1 and 7.1 and you will receive the message to the right if you open the “Room Correction” tab in Quadraphonic or Stereo mode. If 5.1 or 7.1 is installed and setup in the “Speaker Configuration” page, you will see the full Room Correction menu without an error message. All speakers will initially show semi-transparently. To edit this menu, check the box for “Enable Room Correction” at the bottom-center.



Room Correction only works under 5.1 or 7.1 speaker mode.

After you check the box to Enable Room Correction, you will need to select whether you want distances displayed in Feet or Meters. Once selected, the unit of measurement will be colored in light orange, and will be reflected in the measurements below for each speaker. From here, enter the measurements from the seating area to each speaker. Once the distances are entered, test the audio and adjust the gain value to your liking.



The final Device section covers the Microphone. This tab includes both a recording and a playback section.

Both features have similar balance options as the other sections, but differ slightly in the Volume slider function.


The Recording Volume slider pertains to audio recorded by your Microphone.


Raising the volume allows the audio to be recorded at a higher volume; in other words, if the Microphone is used to speak to another person, and the person has a hard time hearing you, raising the volume will make it easier to hear you. Conversely, if you are too loud, lowering the volume slider will produce a quieter audio. Muting this setting will mute all audio recorded by your microphones.

The Playback Volume pertains to a function that allows the Realtek controller to play back any audio the microphones picks up. This is a useful feature to assist with hearing your own speaking volume or audio content playing in the background. By default, this setting is muted. If you encounter an issue where you constantly hear yourself talk, and you're sure it's coming from the computer, check first to see if this setting is enabled.

The microphones Boost button is next to the upper Mute button. The boost assists by increasing the gain of what your microphones picks up while recording. If your microphones is not consistently picking up your voice when speaking, try adjusting the boost higher in small increments, as a little boost can go a long way.

To increase the boost, click the "Microphones Boost" button. When the slider window pops up, adjust the slider bar to the right, then close the popup by clicking the "X" and retest. Whenever you test a different microphones, reset the Boost to 0, test again, and adjust as necessary. You should keep in mind, however, that some mics naturally record louder than others, and too much gain can ruin a recording. Likewise, if you use your microphones for VOIP applications, too much boost can be painfully loud for whomever is on the other side.





Under “microphones Effects,” there are two options that may be enabled: “Noise Suppression” and “Acoustic Echo Cancellation.” “Noise Suppression” helps to remove background noises, such as fans, air conditioning, or anything else that causes a consistent ambient sound in the room.

When enabled, “Acoustic Echo Cancellation” removes echoes caused by room acoustics. For example, when you speak into your mic, your mic initially records your voice, but it may also pick up an echo caused by your voice bouncing off objects in your room. This creates an echo audible to whomever you’re speaking to, much like being on a phone call with someone on speaker phone and you hear your own voice coming back into your speaker. In these situations, this setting can prevent this from occurring.

As with any speaker or microphones setting, your audio equipment will have a significant effect on audio quality and effectiveness of multiple settings. For example, some microphones already equip some degree of noise or echo cancellation and do not benefit further from Realtek’s options. Likewise, even for microphones that do not have built-in noise or echo cancellation, Realtek’s cancellation effects may create distortion or other audio problems. You will need to test your own equipment to determine the effectiveness and usefulness of the various settings outlined above in the Realtek HD Audio Manager guide.


Configuration of the Killer NIC Software

This will give an overview for the installation and the use of the Qualcomm Killer NIC™ Suite. To use the Killer Performance Software, you must first connect a CAT5/CAT5e/CAT6 Ethernet cable to the Killer Ethernet jack (see Page 14 for Killer Ethernet location), and install the Killer Suite software package. This software will NOT work with the Intel NIC.

Please keep in mind that the purpose of the Killer NIC and accompanying software is *NOT* to make your GB/S network suddenly faster; instead, the Killer Suite allows you to configure upstream and downstream bandwidth allocation on an application-by-application basis, giving you a very basic QOS-type set of controls with a simplified GUI. Using the software, you can control which application gets exactly how much upload and download speed/bandwidth. You can also deny bandwidth to applications altogether, as well.

1. Go to www.evga.com/support/drivers, click on the Motherboard tab, select Z270 – Killer NIC – Your OS (Windows 10 x64 was used in this manual), then the submit button at the bottom to download. Once downloaded, extract and run the installer.

Once you have extracted the files, go into the folder labeled “KillerSuite,” and you’ll see files and folders similar to what’s contained in the image below. Double-click on “Setup” to launch the installer.



The Killer NIC installer follows a fairly common software install process. Select your language, click next, ensure that your anti-virus is disabled and click next. Click next on

the Install Shield windows, read and agree to the End User License Agreement, and click next. The initial installation will finish, and then reboot to continue.




After rebooting, the installation will be complete and you may start configuring the Killer Suite to fit your needs. If you're ready to begin, skip down to Page 116 (next page).

If something goes wrong during the installation or you later have network issues, run the driver setup again. This time, instead of a straight-forward installation, the Setup will give you the option to Modify, Repair, or remove the drivers, as seen on the right.

As a troubleshooting step, you may consider first trying the "Repair" option. This option will attempt to discover any missing components or corrupt files in the Killer Suite software and resolve the issue.


If you add a Killer Wi-Fi card or appear to have missing components, the "Modify" screen will give you additional options that may not have installed earlier. There are options that appear to be listed twice, but it is recommended to stick with the default options, which have been auto-detected by the software. After clicking Install, the Setup will finish installing drivers, similar to the screens on the previous page.




Now that the Killer Performance Suite has been installed, you can begin configuring the Suite. You should have both a Desktop and Start Menu icon to launch the Killer Suite application. Double-click on the “K” icon to launch.



Upon launching the Killer Suite you will see this basic UI. The Killer Network Manager is organized into four main sections on the front page of the software suite, giving you full configuration and monitoring options provided by the Killer Suite.



The Application menu allows you to see all applications and processes currently running and/or installed, how much upload/download bandwidth they are currently using, and the ability to control the usage at an application level.




The Applications page will initially be blank the first time it is launched, similar to the above image.

Mousing over the icon in the lower-center will identify which model of Killer NIC is present. If there is no image, then you likely have your CAT5/5e/6 cable plugged into the Intel NIC.


Checking the box for “Show All Applications” will provide you with a larger list of applications and processes to allow for more fine-tuning of your network performance.

Once you have used your Killer NIC for some time, the Application screen will look something like the image to the right. There are 3 sections outlined in red: the checkbox, scroll bar and the bandwidth slider. When the checkbox for “Show all applications” is ticked, you will see many items that you don’t



commonly see in Windows, including background services, inactive processes, and installed but unused applications. This is why the scrollbar on the side is as small as it is, because there will likely be pages worth of information available by the time you scroll to the bottom.


The right side of the program list will have a gray square indicating the percentage of bandwidth this item is permitted to use, but not necessarily *WILL* use. Referring to the pic above, the default is 100%.





The figure above shows the bandwidth controls as well as other information provided on this page. You can adjust allotted bandwidth for applications and processes by setting the Priority of the application/process, which allows the Killer Suite to auto-adjust bandwidth allocations. If you mouse over the “Priority” tab, you will see the path of the

executable that is running. When you click on the pulldown box, you will see the six tiers of priority for the given application. By manually adjusting priorities, you allow the Killer Suite to auto-adjust which programs receive priority bandwidth. It is NOT recommended to change the priority of core Windows processes.


Alternatively, you can manually set a limit for any application or process by using the slider bar. If you look at the images above, you will see the slider box outlined to the right of the application. Drag the slider left or right to manually set the upload and download limits of the application; the slider can be dragged in 1% increments. The overall bandwidth speed for applications is based on the speed you enter in the next section.



The Network section enables\disables bandwidth control, which is the core of the Killer NICs functionality, and also where you manually enter you internet speed.




Without the “Enable Bandwidth Control” option enabled, you cannot utilize the tuning options available in the Killer Suite, including the Applications section above. Once the option is checked, you will get the warning shown in the image above. Because the Killer Suite relies on accurate internet speed data, try an application or website, such as Speedtest.net to find your speed in Megabits (not Megabytes). Although your ISP can offer some details, it is more accurate to provide the Killer Suite with data from an actual test. The more accurate the speed entered, the more effective the Killer Suite will be at controlling bandwidth to various applications and processes.




The next section is Performance, which largely exists to show graphs and other information about individual applications and overall network usage at a macro level. Click on performance and you will see the “System Performance” section, with three main tabs at top: “Top 5 Applications,” “Usage – Download,” and “Usage – Upload.” The fourth tab will be greyed out unless a Killer 1525 or 1535 WiFi card is installed.

Top 5 Applications is a pie chart (see next page) showing overall usage history; the chart will change over time as you use new programs and processes. For troubleshooting purposes, this can also be used to find unknown or unrecognized processes using heavy resources, which may be an indication of a virus, faulty hardware, or Windows issue.




The two (2) usage tabs show a 2-minute window of network activity; all data is displayed in Megabits per second. This will not show what was using the bandwidth, just the total usage over the Killer NIC 2500 (or WiFi 1525/1535 if installed) within the current time frame.



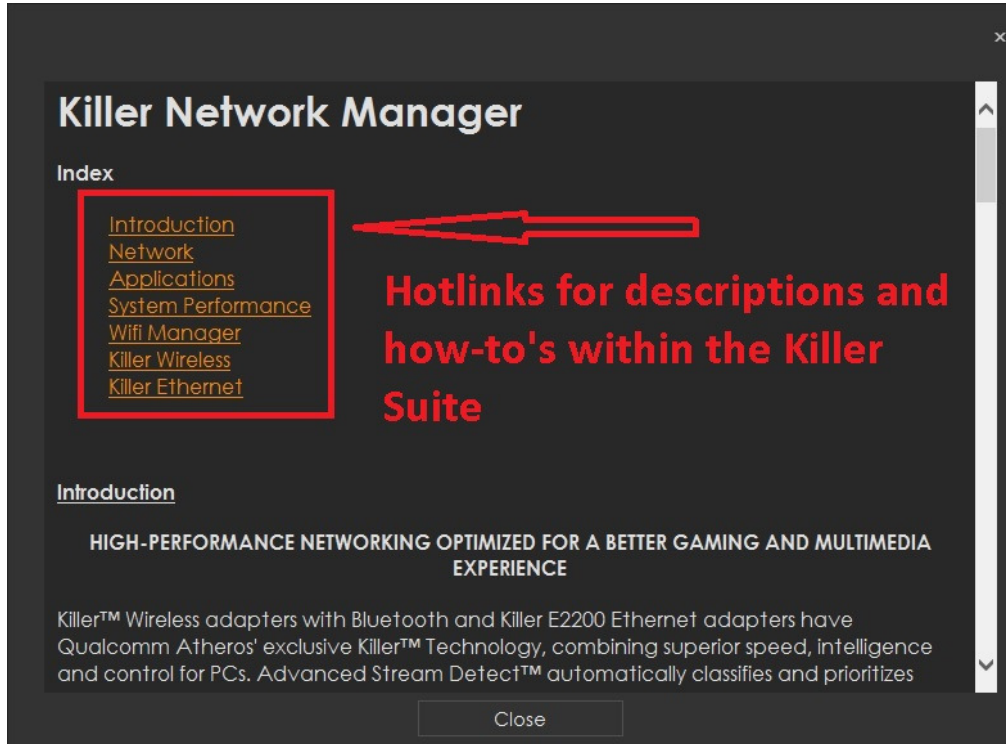
The final section is labeled Killer Ethernet. After clicking Killer Ethernet, this opens to a screen called Network Information if connected to the Killer NIC. If you see an image of a circle with a line crossed through, then you do not have a cable connected to the Killer NIC and the information in this menu will be blank.

If the Killer NIC is connected, you will see information about your Network Information, including speed, IP address, Mac Address, and other information related to your setup.





The Killer Suite also contains a comprehensive help menu. You can enter this help menu by clicking the “?” icon at the top of the any menu. If you click the “?” on the main menu, you will start at the top of the index, and can click any of the hotlinked topics to quickly scroll to the title. Alternatively, if you click the “?” while in any of the four menus, the Help guide will automatically open on the help guide for that section.




This help section within the Killer Suite should be able to cover any addition questions you may still have after reading this overview.


Using the E-LEET Software Suite

EVGA E-LEET is a monitoring and tuning software designed for EVGA motherboards, which is available on the driver DVD and the EVGA website at www.evga.com/E-LEET


After installation, E-LEET will launch directly to the CPU tab and show a general overview of CPU and motherboard information, including Socket information, CPU series, Core speed, core count, cache, etc.





The Memory tab provides an overview of the speed and settings of your system memory. At the top, you can verify that the memory is in Quad/Dual-Channel mode, the size and the NB frequency. Below, you'll find a list of timings for the RAM; not *ALL* of the timings are listed, but the most important and recognizable timings are listed on this tab. For the entire list, you will need to view/adjust those from within the BIOS.



The next tab is “Mainboard,” which provides a bit more detailed information on the motherboard itself. Here, you can find the part number in the “Model” line, the CPU speed, the chipset, and even the date and revision of the BIOS.



The next tab is “Monitoring,” which is an overview of temperatures and voltages in real-time. Please note that all readings on this page are pulled from motherboard sensors, and can change at any time. Temperatures and voltages are measured at fixed intervals, which can often make the readings appear to jump back and forth. Furthermore, voltages on modern motherboards are sometimes directly adaptive based on load; voltage readings **WILL** change based on usage, which is normal. Although E-LEET does offer functionality that will directly affect temperature and voltage readings, that functionality does not exist on this page. To adjust voltages within E-LEET, please see Page 128 for the Voltage tab information.



The next tab is “LED,” which allows adjustments to the onboard LEDs, Heatsink LEDs, and Cover LEDs. To set display behavior, click on the pulldown button under “Mode.” Select the mode you wish to set color(s) for; the change will update and apply immediately.

To set display color, click on the small color box next

to the pulldown button. A color menu will open, similar to the menu found in Precision XOC for selecting and creating colors. Choose a color from the list or create a custom color from the pop-up palette and click “OK”.



Below Voltages is Temperatures. This covers the CPU (socket), PWM, SYS, individual CPU cores, and the CPU package, which reports temperatures from the CPU Register. The bottom section is for Fans, which displays the RPM of all fans currently plugged in to the motherboard using the 3-pin or 4-pin fan header.

The next tab is “Overclocking,” which covers adjustments to the BCLK, individual Cores, and the RING.


The BCLK can be adjusted by either manually moving the slider, or clicking on the arrows to the right. The BCLK can be adjusted by hundredths of a MHz. For example, the slider will show 10478, whereas the monitor will show 104.8 (rounded).

The Cores and the RING are adjusted in the same fashion, except they are adjusted by whole numbers, rather than fractions.

At the bottom is a checkbox for “Turbo Mode,” which is required to adjust clock multipliers. The Restore and Apply Settings buttons are there to either revert to factory default, or make your changed settings live.




The Options tab is for managing your E-LEET profiles. You can save overclock profiles made in the previous section, as well as selecting an option to load a profile at boot. Use care, however, when setting a profile to load on startup; if the overclock is unstable, it may prevent your system from booting to the OS.



Next is the “Processes” tab, which is used to set Processor Affinity. Processor affinity is no longer used as heavily as it was in past hardware and Windows versions. However, this menu can still be used to bypass the inter-load distribution of the CPU and always set a certain CPU core or cores to run a specific process. The processes pulldown list is a snapshot only of the processes in Task Manager when E-LEET launches. Click the “Refresh Processes” button to ensure the processes list is current before making affinity setting changes, or if the process you want is not showing.

To set affinity, click on the pulldown button next to “Process Name.” Select the Process you wish to set affinity for; the file name will appear in the context menu. Finally, select which core or cores you want to assign to this process and click “Set” at the bottom.





The final section is “Voltages,” which provides a level of voltage control similar to what is available in the BIOS. Using the right image, you’ll notice that there are options for multiple voltages, and two options for your CPU Vcore: Adaptive and Override voltages.


Generally speaking, Adaptive will allow you to tailor your CPU Vcore voltage profile to keep voltage usage and heat to a minimum, while Override is more of a brute-force method of setting your CPU Vcore voltage.

Adaptive gives you both Target Voltage and Offset Voltage as an option. When adjusting Adaptive, you start by setting your target voltage, which you will discover through stress-testing your CPU. The Offset voltage is then used to fine-tune additional voltage that may be required in load or stress situations.

Override just provides a static voltage; outside of extreme benching, this setting is typically not needed or ideal for the longevity of the CPU, since the VCore is meant to fluctuate based on load. Alternatively, this can be used in the early stages of overclocking to find your Target voltage, before switching back to Adaptive.



The non-CPU Vcore voltages do not utilize the Adaptive or Override voltage settings. Instead, these use a pulldown menu to select preset voltages, as shown in the image below. The voltages for these components are set by detection during POST, and E-LEET will open the pulldown to the currently detected voltage. Please be careful when adjusting voltages, as there are risks to running electronics out of spec. Although EVGA warranties overclocking, other components are manufactured by different brands (i.e. RAM and CPU), which may have different policies towards overclocking. When your changes are complete, click “Apply Selection” at the bottom, or return to defaults by clicking the “Restore” button.



Below is the voltage range list within E-LEET:

ELEET Voltage Values	Minimum Voltage	Maximum Voltage	Default Voltage
Vcore Target Voltage:	0.800	1.520	-
Vcore Offset Voltage:	-0.500	0.500	0.000
VMesh Target Voltage:	0.800	1.520	-
VMesh Offset Voltage:	-0.500	0.500	0.000
CPU VIN:	1.800	3.050	1.800
DIMM Voltage AB:	1.200	2.000	1.200
DIMM Voltage CD:	1.200	2.000	1.200
PCH 1.0V:	1.000	1.597	1.000
VCCIO:	0.950	1.800	0.950
VSA:	0.700	2.000	1.050

This is the E-LEET utility in its entirety; this should provide enough of an overview to utilize the E-LEET software suite.

Installing Drivers and Software

Note: It is important to remember that before installing the driver CD that is shipped in the kit, you must first load or install your operating system. The motherboard supports 32bit and 64bit versions of Windows 10, 8, and 7.

The kit comes with a CD that contains utilities, drivers, and additional software.

The CD that has been shipped with the EVGA X299 Motherboard contains the following software and drivers:

- ❑ Chipset Drivers
- ❑ Audio Drivers
- ❑ RAID Drivers
- ❑ LAN Drivers
- ❑ USB 3.1 Drivers
- ❑ ME Drivers
- ❑ EVGA E-LEET X
- ❑ User's Manual

Windows 10 Driver Installation

1. Insert the EVGA X299 installation DVD for the motherboard included in the kit.
2. The DVD will auto-run. Install the drivers and utilities listed on the install screen. If the DVD does not run, go to My Computer and click on the CD to open.

Warranty and Overclocking

EVGA's motherboard and graphics card components are designed for and intended for enthusiasts. To be perfectly clear, overclocking will **NOT VOID** your warranty. This means that you are free to do any type or level of overclocking you like, including traditional means, such as by using the BIOS/UEFI (for motherboards); software overclocking with EVGA utilities, such as Precision XOC or E-LEET; or even third-party software utilities. If the motherboard or card fails due to overclocking, you're covered.

Likewise, EVGA understands that enthusiasts are not always satisfied with the provided stock cooling options. Replacing heatsinks, fans, or thermal grease will **NOT VOID** your warranty. Want to install a waterblock on your new motherboard or graphics card? No problem. Planning on a long benching session with a DICE/LN2 pot? Good luck on the benching, but your warranty will still be intact after you've set new world records.

Of course, there are some limitations to our warranties. If an EVGA motherboard or graphics card sustains physical (i.e. damage to the PCB or component due to slippage with a hand tool) or liquid damage, the warranty is void. Similarly, permanent changes to components or the PCB will void your warranty, including the removal of pre-attached components soldered to the board, or the addition of components that are not removable without heat or solvents. In other words, you may use Precision XOC to raise the voltage on your graphics card while overclocking, but you may not solder on another power connector as an alternative.

If the product fails within the warranty period, you must return the product to its original configuration. This means ALL original heatsinks and fans must be reattached as they were when it was first shipped. Even if the product no longer works, it must be returned to its original state before sending it in. If you no longer have the original heatsinks you may be charged for the missing parts.

The EVGA warranty covers overclocking **ONLY** to EVGA products; EVGA is **NOT** responsible for damage caused to components manufactured by other companies. Before overclocking your system, check with the manufacturer of your other components to ensure that you will not void the warranty on those parts due to overclocking. Due to the uniqueness of every system, EVGA Customer Support cannot assist you with overclocking. For overclocking assistance, please visit our Community Forums at <http://forums.evga.com>

Troubleshooting

Replacing a BIOS chip

Replacing a BIOS chip is typically unnecessary. In most cases, this is only needed after one of the following:




- Power failure during a BIOS update
- Failed overclock that does not recover after a hard CMOS reset
- Your motherboard was purchased just as a new line of CPUs were being released and you purchased the motherboard and one of the new CPUs at the same time (e.g. The Intel Broadwell CPUs launched approximately nine months after the Z97 motherboards launched). Chances are, if you purchased the board and CPU together in that timeframe, retail motherboards may not have the latest BIOS with compatibility added for that new CPU, and the board may need a BIOS update to be able to POST successfully.

FOR ANY AND ALL INSTANCES WHERE YOU THINK YOU MAY NEED A REPLACEMENT BIOS CHIP, PLEASE CALL CUSTOMER SERVICE IMMEDIATELY. DO NOT ATTEMPT TO REPLACE THE BIOS CHIP WITH ONE **NOT SUPPLIED BY EVGA, AS EVGA CANNOT OTHERWISE GUARANTEE COMPATIBILITY. CONTACT INFORMATION, HOURS, AND LOCATIONS FOR ALL EVGA CUSTOMER SUPPORT OFFICES CAN BE FOUND HERE:**


<http://www.evga.com/about/contactus/> REMEMBER, ONLY NORTH AMERICAN CUSTOMER SUPPORT IS AVAILABLE 24/7.

The steps for replacement below assume you have already contacted EVGA Customer Support and have received a pre-flashed replacement chip directly from EVGA.

1. Locate the housing for the chip on the motherboard. It should be located on the right side below the RAM, about halfway down the board. The chip will NOT be plainly visible, but under a set of gray plastic doors that keep the chip in place.
2. There is a circular mark on the bottom right that is used for alignment (Circled in RED). Please keep in mind that this chip ****IS DIRECTIONAL**** and inserting the replacement backwards will cause immediate failure of the motherboard and BIOS chip, including the potential for fire or smoke from the socket.
 - a. As a tip, the doors will only close correctly if the chip is inserted correctly. If the chip is in backwards, the door will not fully close without force, which will cause warping or damage to the door.





3. Remove the BIOS chip by lifting it straight up. Set it aside and insert the replacement chip. Once installed, close the left door first and the right door second, making sure both latch into place; you will feel a click when this happens. It is normal if the chip moves a little in the socket when closing the left door. Lastly, close the right door. Once both doors are closed, you've successfully replaced the BIOS chip.



4. **NOTE: THE BELOW IMAGES ARE USED TO SHOW WHAT HAPPENS IF YOU ATTEMPT TO INSTALLED THE BIOS CHIP **INCORRECTLY**:**

- a. If the chip is in backwards, the pins will seat, but the top will overhang the socket slightly, making it impossible to seat the door properly (See picture to the right; the red-circled alignment mark is opposite of where it is supposed to be).
- b. When you push the left door, the chip will lift off the right side pins and the door will **NOT** fully close; even if you attempt to close the right door, it will not close completely. In the last picture, you can see the result of trying to close both doors with the chip installed backwards: neither will fully close. If forced, you will break the doors off.



SSD / HDD is not detected

It is important to note that, as with *ALL* storage devices, if there is a connectivity issue, make sure it is enabled in BIOS. Likewise, if there is a device that shares bandwidth with your SSD or HDD (Page 25), make sure that the desired device is enabled in BIOS, or all other troubleshooting that comes after this section is moot.

If you've physically installed a drive that is not detected in Windows, by Windows setup, or within BIOS, please take the following steps. Component troubleshooting is primarily accomplished through process of elimination testing to deduce a faulty component or a bad connection.

Device is not receiving power:

For devices with external power; such as SATA, U.2, and SATA Express; make sure power is connected to the drive, and that the other side of the cable is attached to the correct port on your PSU. If you are using a mechanical drive (either a platter HDD or an optical drive) you should hear the motor spin up when you first power on the system. Similarly, if you put your hand on the device you should feel the drive spin up, warmth from the motor, and a light on the front from an optical drive. If you see, hear, and feel this, then the device is receiving power, and you may skip to the next section. If you do not see, hear, or feel this, first try the cable on a device that you know works; if the cable works, then the issue is the device in question.

However, if you replace the device and the issue persists, then it is either the PSU, one of its modular cables or the port the cable is plugged into. If you are using a modular cable, try a different cable in the same PSU port, and also test on a different port to isolate the source of the problem. In many cases, the PSU manufacturer may be able to replace the cable if it is bad.

Device has power, but does not show up in the BIOS:

If the device is getting power, then we should expect to see the drive in BIOS, unless you are using a third party controller, such as Marvel; for devices on third-party controllers, go to Page 82. If the drive does not show in the SATA Configuration screen in BIOS, make sure the SATA cable is firmly seated on both the drive side and the motherboard. As noted above, if you are using a device on a connector that shares its bandwidth, you will need to check and/or change the BIOS setting related to your device.

If you've already done this, test a different device using the same cable and SATA port to see if the issue persists. If the issue is not resolved, test a different SATA cable, or a different SATA port. Always be sure to test each step separately, otherwise, if several aspects are changed and drives are now detected, then you will not know what change

actually fixed it. If other devices work on the ports and with the same cables, then the issue is the device in question. If testing shows that either one or more of the ports are bad with a different device, then the motherboard is likely the issue. If one or more of the cables test bad with a different device, then replacing the cable(s) should resolve the issue. These same steps can be used for SATA Express; you can only try a different cable or device with U.2, because there is only one U.2 port.


M.2 devices have a bit less to test due to their connection that provides data and power through a single socket. M.2 devices are installed at an angle, gently pushed down flush with the motherboard, and then screwed into place. M.2 drives only fit one way, and cannot fit upside-down; however, you can have a bad connection if the card is not pushed all the way into the slot before bolting it down. Please refer to Page 32 for M.2 hardware installation, and Page 50 for setup.

Device has power, is seen in BIOS, but cannot be found in Windows:


Lastly, if the device was just added, other troubleshooting in this section has been done, the device has power and shows up in the BIOS, but is not seen in Windows, then it may be that the drive needs to be initialized, partitioned, and formatted in Windows. See Page 82 for instructions.

System does not POST, and POST code indicator reads “C”

When the system powers on, the POST code indicator should cycle through several different codes before booting. However, if the boot process does not complete, you should look at the LED indicator, as it will give you diagnostic information. When a POST code indicator cycles through codes and stops at a specific code, this is an indication that initialization of a specific component or process on the board failed. If you power on the board and it goes directly to “C,” and nothing else, that is an indication of a CPU-related issue. It is not likely a faulty CPU, although it is possible; instead, it is most likely a CPU power issue. Above the CPU socket, there will be one or two 8-pin power connectors (circled in **RED** in the pic to the right). This connector is the main power for the CPU and without power from the PSU, the system cannot fully power on and initialize. Depending on your motherboard, you may have one or two of these connectors; at least one must be plugged in.



Make sure the plug is fully-seated and also make sure that the tab on the socket and the release on the power connector are on the same side (see the box in **RED** on the second picture for the tab on the EPS connector). If the connectors do not line up, then you have plugged in a PCI-E 8-pin into the EPS connector, which may cause irreparable damage to the motherboard or the CPU.



If the PSU is modular, make sure the cable is fully-seated on the PSU side as well. Also, only use cables from your PSU manufacturer for your specific PSU, because using a cable designed for a different PSU may have catastrophic effects. Finally, ensure that the cable on the PSU side is plugged into the connector for the CPU, typically labelled “CPU” or “EPS,” but consult your power supply manual for the correct plug.

System does not POST, and POST code indicator reads “55”

A POST code error reading “55” indicates a memory issue. Verify that memory is installed, and that it is properly seated. Make sure that the memory is on the official support list at www.evga.com/support/motherboard and click on “EVGA X299 FTW-K.” If the memory is not on the list, it may still work because EVGA is unable to test every memory kit released. However, this motherboard will not support modules over 16GB or ECC/Registered RAM.

If your memory should work on this motherboard, the next step is to bring it down to one stick of memory in “DIMM slot 2” and retest. If this is your first boot with the motherboard and the board will not POST, your next step is to press the “Reset CMOS” button. The system will power on when the Reset CMOS button is pressed, which is normal. This should allow the system to boot up. When you see the EVGA X299 FTW-K screen, press “Del” to enter BIOS, set the XMP (instructions are covered in first boot section on Page 48), press F10 to save and exit, which will reboot the system. If the system reboots successfully, you can power down and install the rest of the memory to retest.

System does not POST, and POST code indicator reads “d7”

If the motherboard cannot detect a Keyboard, the system will fail to POST with a “d7” POST code. When this happens, some video card / monitor setups will display a message on screen about a Keyboard error, while others will just remain blank. If this happens, power down the system and reseal the USB connector. Make sure the Keyboard is NOT in the USB 3.1 Type-A connector, as this port requires drivers to initialize and fully function. However, this error can sometimes happen with older wireless keyboards, as well. If this is the case, we recommended that you plug in a regular USB keyboard, reboot, enter BIOS, and go to Advanced – USB Configuration. Make sure “Legacy USB Support” is set to “enabled,” then press F10 to save an exit. Your keyboard should now work, and the board should POST past “d7.”

Have a question not covered above, or want some online resources?

If you have any issues or questions, you can find a lot of valuable help and information on our website, including contact information for our 24/7 US Technical Support/Customer Service toll free at 888-881-3842.

- You may find all of the motherboard drivers at <http://www.evga.com/support/download/>. Select the “Motherboard” tab, then select “Intel X299” from the Family tab, then “142-SX-E297” from the part number pulldown. Next, select either “All” or the specific component that requires a driver, then select your operating system and click “Submit”
- You may find the list of officially supported CPUs and System RAM for your motherboard, including the minimum required BIOS revision for compatibility, at: <http://www.evga.com/support/motherboard/>. BIOS updates are cumulative, so they will always support products from the previous version, except when specified otherwise.
- For our Frequently Asked Questions (FAQ) database, go to: <http://www.evga.com/support/knowledgebase/>
- To visit our Community Forums, go to: <http://forums.evga.com/>
- For our Community Forums section dedicated to X299 motherboards, go to: <http://forums.evga.com/EVGA-X299-Series-f105.aspx>
- Just finished building or upgrading your system and want to show it off? Make a ModsRigs Profile here: <http://www.modsrigs.com/> and create a page dedicated entirely to YOUR system!
- Still building your rig? Make a build log here: <http://forums.evga.com/EVGA-MODS-RIGS-f33.aspx>
- Want to join the online EVGA Gaming Community? Sign up and play with like-minded gamers here: <http://www.evga.com/TEAMEVGA/>

POST Beep codes

POST beeps are used in conjunction with the POST Code indicator to help determine the root cause of the issue when your system is malfunctioning. However, with the advent of modern UEFI/BIOS, the USB detection can throw off the number of beeps making the POST code, while helpful, not as useful as it once was. The motherboard will typically beep 1-2 times per USB device plugged in at boot; some will beep more based on functions of the device. These may occur at the same time as the system's POST beeps, so it can sometimes be a challenge to separate the POST beeps from the USB initialization. Since the POST code indicator now fills most of the role previously held by audible beeps, there are far fewer POST beeps to help debug problems, compared to previous motherboard generations.

Beep Codes	
# of Beeps	Description
1	Invalid Password
4	Some of the Architectural Protocols are not Available
5	No Console Output Device detected
5	No Console Input Device detected
6	Flash Update Failed
7	Reset Protocol is not available
8	Platform PCI resource requirements cannot be met

1 Beep: This relates to BIOS boot passwords. The BIOS will beep, indicating that a password is required to boot. If you incorrectly enter your BIOS password when prompted by the BIOS, you will receive one beep.

4 Beeps: This is an error relating to some small recovery errors, which can occur when coming out of sleep/hibernate.

5 Beeps: This error is stating it cannot detect a keyboard. This can also indicate a video card is not detected.

6 Beeps: This relates to either a failed BIOS update, or a failed Intel ME Firmware update.

7 Beeps: This relates to a failed CPU INIT request. EFI/UEFI Reset protocol is used when a processor INIT call is used and was not successful. This returns either an "EFI_UNSUPPORTED" error if the INIT type is not supported or an "EFI_Device_Error" if the error occurs while generating the INIT; either will result in seven (7) POST beeps.

8 Beeps: Error given when an incompatible PCI Express device is installed and cannot initialize properly.

POST Port Debug LED

Provides two-digit diagnostic POST codes that shows system boot status and can also show why the system may be failing to boot. The LED is extremely useful during troubleshooting situations.

This Debug LED will display a series of hexadecimal (0-F) codes during the POST and will display current CPU socket temperatures after the system has fully booted into the Operating System. See the “POST CODE” section below for more detailed descriptions of specific POST Codes.



Debug LED with CPU
Temperature Monitor

POST Codes

This section provides the AMI POST Codes for the EVGA X299 FTW-K Motherboard during system boot up. The POST Codes are displayed on the Debug LED readout located directly on the motherboard. See Page 12 and 13, component 19 of the Component Legend for physical location.



Debug LED with CPU Temperature Monitor

This Debug LED will switch to display current CPU temperatures after the system has fully booted into the Operating System. Please note that if you see the Debug LED constantly switching numbers after entering Windows, the Debug LED is showing your CPU temperature, and is NOT reporting errors.

AMI POST Codes

1	Power on. Reset type detection (soft/hard).
2	AP initialization before microcode loading
3	North Bridge initialization before microcode loading
4	South Bridge initialization before microcode loading
5	OEM initialization before microcode loading
6	Microcode loading
7	AP initialization after microcode loading
8	North Bridge initialization after microcode loading
9	South Bridge initialization after microcode loading
0A	OEM initialization after microcode loading
0B	Cache initialization
0C-0D	Reserved for future AMI SEC error codes
0E	Microcode not found

0F	Microcode not loaded
10	PEI Core is started
11-14	Pre-memory CPU initialization is started
15-18	Pre-memory North Bridge initialization is started
19-1C	Pre-memory South Bridge initialization is started
1D-2A	OEM pre-memory initialization codes
2B	Memory initialization. Serial Presence Detect (SPD) data reading
2C	Memory initialization. Memory presence detection
2D	Memory initialization. Programming memory timing information
2E	Memory initialization. Configuring memory
2F	Memory initialization (other).
30	Reserved for ASL (see ASL Status Codes section below)
31	Memory Installed
32	CPU post-memory initialization is started
33	CPU post-memory initialization. Cache initialization
34	CPU post-memory initialization. Application Processor(s) (AP) initialization
35	CPU post-memory initialization. Boot Strap Processor (BSP) selection
36	CPU post-memory initialization. System Management Mode (SMM) initialization
37-3A	Post-Memory North Bridge initialization is started
3B-3E	Post-Memory South Bridge initialization is started
3F-4E	OEM post memory initialization codes
4F	DXE IPL is started
50	Memory initialization error. Invalid memory type or incompatible memory speed
51	Memory initialization error. SPD reading has failed
52	Memory initialization error. Invalid memory size or memory modules do not match.
53	Memory initialization error. No usable memory detected
54	Unspecified memory initialization error.
55	Memory not installed
56	Invalid CPU type or Speed
57	CPU mismatch
58	CPU self-test failed or possible CPU cache error
59	CPU micro-code is not found or micro-code update is failed
5A	Internal CPU error

5B	reset PPI is not available
5C-5F	Reserved for future AML error codes
E0	S3 Resume is started (S3 Resume PPI is called by the DXE IPL)
E1	S3 Boot Script execution
E2	Video repost
E3	OS S3 wake vector call
E4-E7	Reserved for future AML progress codes
E8-EB	S3 Resume Failed
EC-EF	Reserved for future AML error codes
F0	Recovery condition triggered by firmware (Auto recovery)
F1	Recovery condition triggered by user (Forced recovery)
F2	Recovery process started
F3	Recovery firmware image is found
F4	Recovery firmware image is loaded
F5-F7	Reserved for future AML progress codes
F8	Recovery PPI is not available
F9	Recovery capsule is not found
FA	Invalid recovery capsule
FB-FF	Reserved for future AML error codes
60	DXE Core is started
61	NVRAM initialization
62	Installation of the South Bridge Runtime Services
63-67	CPU DXE initialization is started
68	PCI host bridge initialization
69	North Bridge DXE initialization is started
6A	North Bridge DXE SMM initialization is started
6B-6F	North Bridge DXE initialization (North Bridge module specific)
70	South Bridge DXE initialization is started
71	South Bridge DXE SMM initialization is started
72	South Bridge devices initialization
73-77	South Bridge DXE Initialization (South Bridge module specific)
78	ACPI module initialization

79	CSM initialization
7A–7F	Reserved for future AMI DXE codes
80–8F	OEM DXE initialization codes
90	Boot Device Selection (BDS) phase is started
91	Driver connecting is started
92	PCI Bus initialization is started
93	PCI Bus Hot Plug Controller Initialization
94	PCI Bus Enumeration
95	PCI Bus Request Resources
96	PCI Bus Assign Resources
97	Console Output devices connect
98	Console input devices connect
99	Super IO Initialization
9A	USB initialization is started
9B	USB Reset
9C	USB Detect
9D	USB Enable
9E–9F	Reserved for future AMI codes
A0	IDE initialization is started
A1	IDE Reset
A2	IDE Detect
A3	IDE Enable
A4	SCSI initialization is started
A5	SCSI Reset
A6	SCSI Detect
A7	SCSI Enable
A8	Setup Verifying Password
A9	Start of Setup
AA	Reserved for ASL (see ASL Status Codes section below)
AB	Setup Input Wait
AC	Reserved for ASL (see ASL Status Codes section below)
AD	Ready To Boot event

AE	Legacy Boot event
AF	Exit Boot Services event
B0	CPU Memory controller configuration
B1	Runtime Set Virtual Address MAP End
B2	iMC init
B3	Memory training
B4	Memory training
B5	Memory training / timing training
B6	Memory training
B7	Memory training
B8-BF	Memory training / DRAM final configuration
C0-CF	OEM BDS initialization codes
D0	CPU initialization error
D1	North Bridge initialization error
D2	South Bridge initialization error
D3	Some of the Architectural Protocols are not available
D4	PCI resource allocation error. Out of Resources
D5	No Space for Legacy Option ROM
D6	No Console Output Devices are found
D7	No Console Input Devices are found
D8	Invalid password
D9	Error loading Boot Option (Load Image returned error)
DA	Boot Option is failed (Start Image returned error)

EVGA Glossary of Terms

AC – Alternating Current

ACPI - Advanced Configuration and Power Interface

AHCI – Advanced Host Controller Interface

APIC - Advanced Programmable Interrupt Controller

BCLK – Base Clock (or operating frequency of base system bus)

BIOS - Basic Input Output System

CMOS - Complementary Metal-Oxide Semiconductor

CPU – Central Processing Unit

DC – Direct Current

DDR - Double Data Rate

DIMM - Dual In-line Memory Module

DMI – Direct Memory Interface

DP – Display Port

DRAM - Dynamic random access memory

DVI – Digital Video Interface

E-LEET/E-LEET X – EVGA motherboard monitoring and tuning software

FIVR – Fully Integrated Voltage Regulator

GHz – Gigahertz

GPIO (Thunderbolt) – General Purpose Input/Output

GPU – Graphics Processing Unit

GUI – Graphical User Interface

HDD - Hard Disk Drive

HDMI - High-Definition Multimedia Interface

HDR – High Dynamic Range (Lighting)

HPET - High Precision Event Timer

HT – Hyper-Threading

HSF - Heat Sink Fan

I/O - Input/Output

IEEE - Institute of Electrical and Electronics Engineers

IGP - Integrated Graphics Processors

IMC – Integrated memory controller
IOH – Input/Output Hub
IRQ - Interrupt Request
JBOD - Just a Bunch of Disks
JEDEC - Joint Electron Device Engineering Council
LAN - Local Area Network
LCD - Liquid Crystal Display
LGA – Land Grid Array
LN2 – Liquid Nitrogen Cooling
MAC - Media Access Control
M.2 – A card bus based standard commonly used for VERY fast SSD's, and WiFi.
MCP - Media and Communications Processor
Intel ME – Intel Management Engine
MHz – Megahertz
MMIO – Memory Mapped I/O
NB - Northbridge
NCQ - Native Command Queuing
NIC - Network Interface Card
NTFS - New Technology File System
NVMe – Non-Volatile Memory Express
OEM - Original Equipment Manufacturer
PATA - Parallel Advanced Technology Attachment
PCB - Printed Circuit Board
PCH – Platform Controller Hub
PCI - Peripheral Component Interconnect
PCI-E - Peripheral Component Interconnect Express
PLL – Phase Locked Loop
POST – Power on Self-Test
PWM – Pulse Width Modulation
QDR - Quad Data Rate
QOS – Quality of Service

QPI – Quick Path Interconnect
RAID - Redundant Array of Inexpensive Disks
RAM – Random Access Memory
ROM – Read Only Memory
RGB - Red Green Blue
SATA - Serial Advanced Technology Attachment
SAS – Serial Attached SCSI
SB - Southbridge
SCSI - Small Computer System Interface
SFR – Split Frame Rendering
SLI - Scalable Link Interface
SPD - Serial Presence Detect
S/PDIF - Sony/Philips Digital Interconnect Format
SPP - System Platform Processors
SSD – Solid State Drive
TCP/IP - Transmission Control Protocol/Internet Protocol
Thunderbolt – High-speed hot-swappable interface (similar function to USB)
U.2 – High-Speed storage interface previously referred to as SFF-8639
UDP – User Datagram Protocol
UEFI – Unified Extensible Firmware Interface
USB - Universal Serial Bus
VDroop - Voltage Droop
VGA - Video Graphics Array
VREG – Voltage Regulator
VROC – Virtual RAID On CPU

Compliance Information

FCC Compliance Information

This device complies with FCC Rules Part 15. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation. This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the manufacturer's instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures: (1) Increase the separation between the equipment and signal source, or (2) connect the equipment to an outlet on a circuit different from that to which the signal source is connected. Consult the dealer or an experienced computer technician for help. The use of shielded cables for connection of peripheral devices to the PC systems is required to ensure compliance with FCC regulations. Changes or modifications to this unit not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.

CE Compliance Information

Generic Radiation Interference Standard for Information Technology Equipment. (EN 55022: 2006, Class B), (EN 61000-3-2: 2006), (EN 61000-3-3: 1995 + A1: 2001 + A2: 2005). Warning: This is a Class B product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measure. Generic Immunity Standard for Information Technology Equipment. (EN 55024: 1998 + A1: 2001 + A2: 2003).

Trademark & Copyright Information

2001-2017 EVGA Corp. EVGA, the EVGA logo and combinations thereof are trademarks of EVGA Corp. All brand names, company names, service marks, logos, and trademarks of the company, or its affiliates or licensors are trademarks or registered trademarks of the company or its subsidiaries, affiliates or licensors in the US and other countries. Other company, products and service names may be trademarks or service marks of others. EVGA reserves the right to terminate this license if there is a violation of its terms or default by the Original Purchaser. Upon termination, for any reason, all copies of Software and materials must be immediately returned to EVGA and the Original Purchaser shall be liable to EVGA.com CORP for any and all damages suffered as a result of the violation or default.

Legal Information

All material including but not limited to, text, data, design specifications, diagnostics, graphics, logos, reference boards, files, images, drawings, and software including this document and the software itself (together and separately) is owned, controlled by, licensed to, or used with permission by EVGA Corporation and is protected by copyright, trademark, and other intellectual property rights. All is being provided "as is", EVGA Corporation makes no warranties, whether express or implied, statutory or otherwise with respect to the materials and expressly disclaims all implied warranties of non-infringement, merchantability, and fitness for a particular purpose. In no event shall the liability of EVGA Corporation for claims arising from the use of the materials by anyone exceed the original purchase price of the materials (or replacement of the materials at EVGA Corporation's option). All information furnished is believed to be accurate and reliable. However, EVGA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use, or use of the Software. No license is granted by implication or otherwise under any patent or patent rights of EVGA Corporation except as expressly provided herein. All specifications mentioned in this publication are subject to change without notice.

Ver. 2