

## Configuring Accelerated-X Wall Display

with the  
Summit HX Series

Some "X" Definitions:

**Display** - An X server can produce one, and only one, "display." If a computer system has only one X server (the normal case), the display is designated "display 0" by default. The display may be very complicated, employing many graphics chips, many monitors, many xscreens - each with different color depths, image overlays, video windows, and so on, or the display be relative simple, with one xscreen using one graphics chip driving one monitor, showing a 2D image without overlays or video windows. A display will usually have a mouse and a keyboard as operator as interactive input devices.

**Cursor** - A display has one cursor, usually shown as an arrow, but can have different shapes and colors, and its position is normally controlled by the operator's mouse. The HX Series X servers allow the territory that a cursor is permitted to travel to be configured, including "wrap-around" parameters if desired. The cursor is not considered a part of the image over which it moves. (This is important when recording for later playback).

**Xscreen** - An xscreen is a fundamental element of a display, and can be relative simple, or rather complicated. Objects in an xscreen are referenced by pixel location from the xscreen's top, left x,y pixel origin - 0,0. Xscreens are numbered 0,1,2, ... n. xscreens have a rectangular shape. Each xscreen has its own set of parameters such as x,y pixel dimensions, color depth, overlays or not, etc. The HX Series X servers allow an xscreen to be assigned to one or more monitors. Conversely, one or more (up to four) xscreens can be assigned to one graphics chip, depending upon the number of monitors a graphics chip is supporting. Windows created by graphics applications can be moved (or dragged) around within an xscreen, but not dragged from one xscreen to another (unlike the cursor, which is allowed to move from xscreen to xscreen).

**Monitor Display Devices** - A monitor is a device that can produce a "view" of one xscreen or a piece of an xscreen. Most modern graphics cards can drive two or more monitors from a single graphics chip by utilizing multiple output paths on a card, each terminating in a monitor connector on the back of the card (sometimes referred to as an "output channel").

**View** - A monitor shows a "view" rather than a "screen." This terminology is required since what a monitor is showing is no longer guaranteed to be an xscreen (as in "the old days" when a graphics card had one graphics chip that could drive one monitor and the image on the monitor was one (complete) xscreen. With SingleLogicalScreen, one monitor may have a small portion of an xscreen as the "view" that it is showing. For example, a Wall Display with one (big) SLS and a dozen monitors, each monitor view is only one-twelfth of the single xscreen in the display..

**Multiple Xscreens per Graphics Chip** - With HX Series X servers, it is usually possible to use an xscreen per graphics chip output channel, if desired. Thus with Matrox QiD cards, for example, HX Series X server can employ one, two, three, or four xscreens per QiD graphics chip, and assign each xscreen to a separate monitor. Often a monitor is referred to as a "head," or as a "screen." Head is correct usage, but screen is no longer correct, as explained in the previous paragraph.

**Pixel Size** - Until a display is being realized on monitor hardware, the actual pixel size is not a concern. However, when building Wall Display Systems with non-identical monitors, the actual physical size of a pixel on each monitor may be important if an xscreen employs more than one monitor, so the "dots per inch," or pixels per inch shown by the monitors grouped together in an xscreen is a parameter that the system designer needs to consider.

**Single Logical Screen** - Also known as "SLS," "Stretched Desktop," and "SingleLogicalScreen" (within Xi Graphics literature), these terms refer to the ability to have an xscreen that employs or uses multiple monitors and/or graphics chips to produce a display or part of a display, such as a large Wall Display System. The Accelerated-X Summit HX Series X servers incorporates the SingleLogicalScreen feature, but it also allows multiple xscreens that employ more than one monitor per xscreen as well as more than one graphics chip per xscreen. And each xscreen is fully hardware accelerated. This ability to mix multiple SLS and multiple non-SLS xscreens, all of which are hardware accelerated, in a single display is unmatched in the graphics display industry.

**Rotated Monitors** - Some monitors have the ability to be physically rotated ninety degrees to the "portrait mode" position. Often, other than the rotation of the physical hardware, nothing within the electronics is changed. Thus the graphics pixel data to be output to a rotated monitor must also be "rotated" to cause the scan lines - which now are vertical - to draw the image correctly. This is a non-trivial piece of work for an X server, especially if multiple monitors and/or graphics chips are involved as in large Wall Displays with one or more SLS xscreens. With the HX Series, this rotation is accomplished efficiently, causing only a minor performance (speed) hit. This is true even if there are multiple SLS xscreens operating in a single display system, and each xscreen has one or more rotated monitors.

**Video In** - The HX Series can support up to six synchronous video inputs (not to be confused with asynchronous video streams that require video capture operations), each being shown in a video window on one or more xscreens. Note that a video window can be moved about in an xscreen, but cannot be dragged to another xscreen. Video windows in HX Series can be scaled, and can be moved across monitor and graphics chip boundaries as long as the window does not leave its assigned xscreen. So if a large Wall Display had a dozen monitors being driven by six, say, graphics chips on three graphics cards and was configured as one SingleLogicalScreen covering all dozen monitors, the video windows can be scaled to be rather large, and can also be dragged about freely.

**Video Capture** - Live incoming video streams are asynchronous to the X server operations, and hence must be synchronized. This is accomplished by "frame grabbing" hardware that is usually - but not always - separate from the graphics chips on a card. Once frames have been captured and buffered, the X server can deal with the video data much in the same way as recorded video from MPEG devices. Currently the HX Series does not have frame grabbing capabilities, but it can be added as a custom project.

**Single-head Cards** - Graphics cards that have one chip and can drive only one monitor are usually referred to as single-head cards. These cards can support only one xscreen.

**Dual-head Cards** - A graphics card that has two graphics chips, each driving a single monitor, or a graphics card with a single graphics chip driving two monitor is referred to as a dual-head card, and such cards can support two xscreens, or two views.

**Triple-head Cards** - Matrox makes some graphics cards with a single graphics chip which can support three monitors. These are triple-head cards, and can support one, two or three xscreens or three views.

**Quad-head Cards** - Graphics cards that can drive four monitors are known as quad-head cards. These come in three hardware configurations: a card with one graphics chip driving four monitors, a card with two graphics chips each driving two monitors, and a card with four graphics chips each driving one monitor. In all cases, with the HX Series, each card can have up to four xscreens. A card with only one graphics chip that is driving four monitors is obviously sharing the "horsepower" of the single chip among the four monitors, and possibly four xscreens, whereas a graphics chip with four graphics chips and four monitors has four graphics engines to support four monitors and up to four xscreens.

**"Free SLS"** - A graphics card that has one graphics chip supporting two or more monitors can be configured by Summit Series software to effect SLS capability easily. Since there is essentially a single frame buffer, Summit's DualView or TripleView or QuadView feature (depending upon the number of monitors being supported) allows each monitor to show "its piece" of the frame buffer image. This is basically free SLS, since no chip (frame buffer) boundaries are crossed, making the SingleLogicalScreen effect simple, easy, and fast. Three monitor SLS OpenGL, anyone? (This is available in the Summit WX Series, which supports OpenGL, but not in HX Series, which does not)

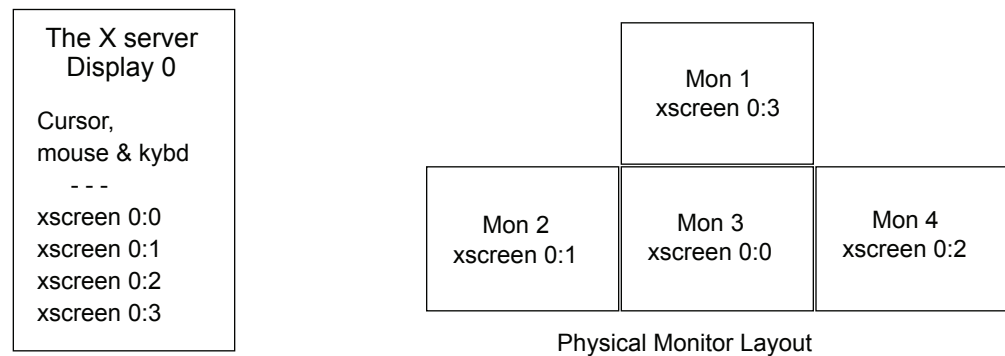
**Computer Platform** - Summit HX Series - as well as other Summit Series - run on x86 (both 32- and 64-bit), and SPARC computers, or "platforms" under UNIX operating systems. SPARC platforms use Solaris for SPARC, while Linux or Solaris for x86 are usually used for x86 platforms.

## The Configuration Process

From the foregoing, it is obvious that the HX Series is very powerful and flexible. The flexibility is almost endless, and thus an understanding of the configuration process and an organized approach to it is helpful when configuring a large and/or complicated display, such as one a dozen or more monitors, more than one "stretched desktop" mixed with some single-monitor xscreens, video and image overlays, etc.

### Example 1. Simple "Inverted T"

Assume a graphics card with four graphics chips, each chip driving one monitor (Matrox's G450MMS, for example), utilizing one xscreen per chip. Overlays are not being used. The X display for this example consists of the cursor, mouse and keyboard, and four xscreens using 24 bit color for 2D displays. The system is to have four identical monitors (for ease of the discussion) physically arranged with three on the bottom and one on the top (centered)



The final assignment of which xscreen is assigned to which monitor is strictly arbitrary, according to the system designer/builder wishes, so let's assume that for some strange reason we wish to have the monitors assigned to xscreens as shown in the monitor (geometry) layout.

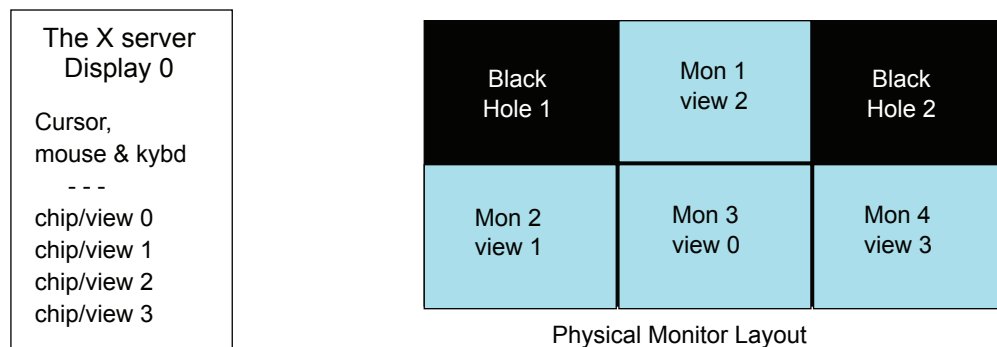
To start the process, the operating system is booted, a step that will cause the OS to locate a "graphics device" (graphics chip or graphics engine) that will be used as the boot device for the console display for the boot process. If all four monitors are cabled to the graphics card in this system, one of them will be used as the OS console display. With the Matrox G450MMS, the boot ROM on the card will usually designate the top connector and its associated graphics chip to be the boot (VGA) path to the monitor that will be the console display, and the the first xscreen, xscreen 0:0, will be assigned to this graphics chip by the X server when it is initially started and assumes control of the console. Other cards may exhibit other characteristics, but which ever monitor "lights up" is the boot console monitor, and will be, by default, xscreen 0:0's "view." If not already installed, now install the Summit HX Series demo package that includes the X server, xsvc kernel module and G450 graphics driver.

Keeping things simple initially, invoke 'xsetup' while in the boot text console, and configure the X server (which is not yet running) for everything except the multiple xscreens and multiple monitors - i.e., the graphics context stuff such as color depth, monitor resolution, and so on. When the basic single xscreen and single monitor setup is configured, start the X server with 'xaccel' and

see if it does what you want with the one xscreen/monitor setup. If it does, kill the X server and restart it with 'startx' to also get the window manager and GUI up with the X server. Now invoke 'xsetup' again which will bring up the Accelerated-X graphical install utility which makes it easy to add to the display the three additional monitors (by specifying "quad") and three more xscreens by clicking three times on the "add xscreen" button. Move the monitor boxes around to get the "inverted T" geometry desired. The xscreens may not be configured as desired, so just drag them to the box they should be in (dragging an xscreen into a box already occupied is not allowed, so empty a box by dragging its xscreen to one of the "empty holes" temporarily to get it out of the way). Check to see that the geometry is correct and then save it, get out of the session and restart the X server again with 'startx.' Done. Simple process for a simple configuration.

### Example 2. "Inverted T" with One Xscreen (SLS)

Assume a the same inverted T monitor layout as in the first example, except in this case the display has one SingleLogicalScreen being generated by the four graphics chips each driving a monitor showing showing a "view" associated with a graphics chip/monitor pair. The term "view" is used in place of xscreen in example 1, since in this example, there is only one xscreen in



display 0 being generated by the four chip/monitor pairs. Windows can be dragged from monitor to monitor, since all monitors are showing only a part of the single xscreen, whereas in Example 1 each monitor was showing a complete xscreen. Since an xscreen must be rectangular, there are two empty spaces each the size of a monitor in this SLS layout. These are referred to as "black holes," since a window or a cursor being moved into one of these spaces will disappear. They are "there" but cannot be seen. Since many GUIs and applications have a tendency to make use of the upper left corner of a display to park items, the inverted T can cause some grief if the applications are not designed for it in an SLS configuration.

The configuration process is similar to Example 1, except there is only one xscreen, and views instead of xscreens being assigned to monitors. Cursor control is discussed in a different section of this paper. Again, configuration is a breeze.

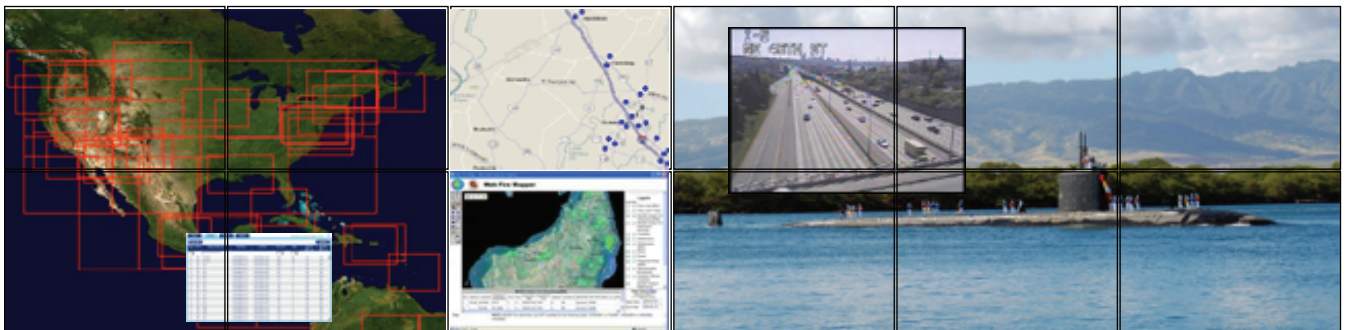
### Example 3. Another Inverted "T" Layout

This example is similar to the previous one but with the top monitor rotated and showing its own xscreen. The three bottom monitors are using one SLS. (Photos are used in the illustrations instead of computer generated graphics due to of laziness of the author). This configuration can be accomplished with two Matrox cards, a DL256PCI, and a P750 AGP or a PCI version of a QID if there is no AGP slot available in the computer .



### Example 4. Twelve-monitor Display w/Video and Image Overlays

Moving right along, we will now configure a rather larger Wall Display that incorporates a dozen monitors in a 6-wide, 2-high configuration, with a couple of SLS areas, some single-monitor xscreens, and both image and video "overlays" on main 24-bit image planes. Keep in mind that all of the display is fully hardware accelerated to the capability of the underlying graphics hardware being used. This implies that if image overlays are required, graphics chips that have image overlay plane support will be highly desirable to keep the graphics performance up. For simplicity sake, this example does not employ any rotated monitors. The width of the display is limited to 32,768 pixels; hence very large monitor resolutions can be handled by the SW.



This display has two "stretched" xscreens and two single-monitor xscreens. All monitor views are fully hardware accelerated. Note that the two video windows - one in each SLS - can be moved about within their respective SLS.