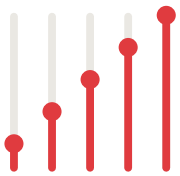

MEASURING
BROADBAND
AMERICA

Technical Methodology
Data Usage

Date:	4 September 2018
Document owner:	SamKnows
Subject:	Data usage test description

Data Usage

This test passively measures the amount of data used on the user's broadband connection. This test does not inspect or record the contents of packets; it only records volumes. Volumes are reported hourly and are broken down by customer wired traffic, customer Wi-Fi traffic and SamKnows test traffic (i.e. the volume of traffic used in performing all of the SamKnows measurements). All volumes are reported independently (one value is not inclusive of any other) and are captured at the MAC level, meaning that they include Ethernet, IP and TCP/UDP header overheads.

The precise mechanism used to record data usage volumes varies by the hardware platform in use. In some hardware platforms a managed hardware Ethernet switch is present. A managed hardware Ethernet switch means that there is a separate processor dedicated entirely to switching Ethernet traffic, which in turn means that the main CPU does not have the overhead of performing Ethernet switching. In these cases the port counters from the managed switch are simply interrogated and recorded (the CPU does not even see the traffic). In other devices only a software switch is present, and in these cases the counters from the software bridge in the Linux kernel are taken. For Wi-Fi, if the hardware platform in question does not run a wireless access point then wireless traffic is monitored passively, and the volumes of the packets will be recorded (regardless of whether they are encrypted or not). If the hardware platform runs a wireless access point, then the wireless traffic will be visible locally in unencrypted form, and the volumes will simply be recorded from the network interface's counters.

The measurement client records the following each hour:

- SamKnows test traffic received
- SamKnows test traffic transmitted
- Customer wired LAN traffic received
- Customer wired LAN traffic transmitted
- Customer Wi-Fi traffic received
- Customer Wi-Fi traffic transmitted

Support by Whitebox model

The following table summarises the approaches used to capture data usage for each model of Whitebox:

Whitebox model	LAN capability	LAN detection approach	WLAN capability	WLAN detection approach
SKWB8	5x 1Gbps	Hardware (SwLib)	2.4Ghz & 5Ghz radios, 802.11a/b/g/n/ac	Passive monitor mode (libpcap)
AC1750v2	5x 1Gbps	Hardware (SwLib)	2.4Ghz & 5Ghz radios, 802.11a/b/g/n/ac	Passive monitor mode (libpcap)
WDR3600	5x 1Gbps	Hardware (SwLib)	2.4Ghz & 5Ghz radios, 802.11a/b/g/n	Passive monitor mode (libpcap)
WR1043ND	5x 1Gbps	Hardware (SwLib)	2.4Ghz radio, 802.11b/g/n	Passive monitor mode (libpcap)
WR741ND	5x 100Mbps	Software (SysClass)	2.4Ghz radio, 802.11b/g/n	Passive monitor mode (libpcap)
WR741NDv4	5x 100Mbps	Software (SysClass)	2.4Ghz radio, 802.11b/g/n	Passive monitor mode (libpcap)
WNR3500L	5x 1Gbps	Software (SysClass)	2.4Ghz radio, 802.11b/g/n	Hardware (BCM wl)

Please note that the LAN capabilities of the Whiteboxes detailed above denote their switching capabilities of local LAN traffic; It does not reflect the maximum speed that can be measured over the WAN Interface.

The data usage approaches referenced above are as follows:

Hardware - SwLib (Wired LAN)

Traffic counters are obtained from a managed hardware switch inside the Whitebox. The user's traffic does not pass through the main CPU, as the hardware switch performs all switching using an onboard ASIC processor. In this scenario, the switch ASIC exposes a management interface that allows user-space applications to configure the switch and retrieve statistics from it. In OpenWrt, the operating system used on Whiteboxes, this switch configuration utility is called "swconfig" and is available in library form as well ("swlib"). This is documented at https://wiki.openwrt.org/doc/techref/swconfig

Software - SysClass (Wired LAN)

On devices that do not have a hardware switch, the bridging of wired LAN to WAN traffic is performed in software using the bridging model of the Linux kernel. In this instance, the counters from the /sys/class/net interface are used. This is a standard interface exposed by the Linux kernel for capturing network interface counters, and is documented at https://www.kernel.org/doc/Documentation/ABI/testing/sysfs-class-net.

Passive monitor mode – libpcap (Wireless LAN)

The majority of Whiteboxes do not run a wireless access point and instead passively monitor wireless activity of the user's wireless network(s) - i.e. they operate as a passive wireless station. The passive monitor mode approach described below is used in such environments.

In this scenario, the Linux utility “iwconfig” is used to place the wireless interface(s) into monitor mode on the appropriate channels. The library “libpcap” is used to capture packet headers. Libpcap captures are configured with BPF (Berkeley Packet Filter) filters that capture traffic to/from the user’s wireless access point only. Additional filters are setup to capture only traffic that is destined to be exchanged with the 802.11 ‘distribution system’. The Rx bytes filter captures packets where the 802.11 fromDS field is set to true and address 2 matches the selected access point's MAC address. The Tx bytes filter captures packets where the 802.11 toDS field is set to true and address 1 matches the selected access point's MAC address. Please note that the use of the fromDS and toDS field captures not only wireless traffic destined to/from the internet, but also to/from wired LAN devices too.

Hardware – BCM wl counter (Wireless LAN)

The WNR3500L Whitebox is unlike all other Whitebox models. It runs a wireless access point that users are instructed to connect their wireless devices to. This is because the WNR3500L was originally intended to replace the user’s CPE, and (by default) operates as a router with a fully featured administrative user interface. Note that it can also operate as a bridge, but this does not affect the wireless operation whatsoever, it still runs a wireless SSID in this configuration.
By virtue of the WNR3500L running its own wireless SSID, there is no need to passively scan for other wireless networks and monitor them. It is only necessary to monitor the internal wireless counters visible within the device. The WNR3500L uses a Broadcom wireless LAN chipset, and Broadcom provide a proprietary wireless LAN driver and management utility named “wl”. The wl utility allows wireless LAN counters to be retrieved directly from the wireless chipset.

The “wl wme_counters” output gives separate counters for the four WME access categories - voice, video, best effort and background. The best effort category (AC_BE) counters are used as wireless traffic counters. The other categories show negligible counts and are excluded.

Handling counter overflows

Whiteboxes often run for extended periods of times without rebooting, so the raw counters obtained by the different counter approaches can be quite large and can roll over as well (e.g. they exceed the limits of a 32-bit signed integer). To minimise the complexity in handling overflows when adding large numbers and in detecting rolled over counters, all counter types first obtain the deltas between the current and previous values of all raw counters. All counter arithmetic calculations are limited to two operands and overflows are handled for each calculation.

Data dictionary

Field	Type	Notes
unit_id	Integer	The unique identifier of the Whitebox
dtime	Integer	Time report printed in UTC (all other fields refer to the preceding 60 mins)
sk_tx_bytes	Integer	Bytes transmitted by SamKnows tests in the report interval
sk_rx_bytes	Integer	Bytes received by SamKnows tests in the report interval
cust_wired_tx_bytes	Integer	Bytes transmitted by customer's wired LAN devices in the report interval
cust_wired_rx_bytes	Integer	Bytes received by customer's wired LAN devices in the report interval
cust_wifi_tx_bytes	Integer	Bytes transmitted by customer's wired LAN devices in the report interval
cust_wifi_rx_bytes	Integer	Bytes received by customer's wired LAN devices in the report interval
location_id	Integer	An internal field used by SamKnows, this can be ignored

#END#