



sandvine[®]

Intelligent Broadband Networks

Global Internet Phenomena Spotlight

Africa, Fixed Access, Fall 2011



Contents

Executive Summary	3
Average Day Network Overview	4
Data Consumption Characteristics	8
Focus on Peak Period	9
Study Details	12
Explanation of Categories	13
Bibliography	14

Figures

Figure 1 - Average Day (Subscribers and Traffic) - Africa, Fixed Access	4
Figure 2 - Average Day (Upstream and Downstream) - Africa, Fixed Access	5
Figure 3 - Average Day (Downstream-to-Upstream Ratio) - Africa, Fixed Access	5
Figure 4 - Downstream-to-Upstream Ratio by Category (24 Hour) - Africa, Fixed Access	6
Figure 5 - Average Day (Network Upstream) - Africa, Fixed Access	6
Figure 6 - Average Day (Network Downstream) - Africa, Fixed Access	7
Figure 7 - Percentile Byte Contribution (Monthly) - Africa, Fixed Access	9
Figure 8 - Peak Period Traffic Composition - Africa, Fixed Access	10
Figure 9 - Prime Time Ratio (Aggregate Traffic) - Africa, Fixed Access	12

Tables

Table 1 - Median Consumption - Africa, Fixed Access	8
Table 2 - Mean Consumption - Africa, Fixed Access	8
Table 3 - Mean-to-Median Ratio - Africa, Fixed Access	8
Table 4 - Top Peak Period Applications by Bytes - Africa, Fixed Access	11



Executive Summary

Africa is seen by many as the next great broadband market, but it is still very much emerging. Nevertheless, the future is very bright - in most of the continent, roughly half of the population is under 19 years of age, and consequently has relatively little purchasing power. As this group grows in wealth, broadband market penetration can increase.

Additionally, Africa is presently served by relatively few major international cables and connection points, but in the coming years that will change as more, and bigger, undersea cables come online. Greater international connectivity means that Africa's traffic profiles will rapidly change as speeds increase and content becomes more readily available. The result is that the traffic profiles seen today very much represent the "before" in a "before and after" comparison.

Examination of Internet traffic on fixed access networks in Africa leads to a number of broad trends and conclusions.

Averaged over 24 hours, P2P Filesharing is the dominant upstream application category, accounting for 49.4% of bytes. With relatively few applications making extensive use of upstream capacity, this level of P2P Filesharing is not uncommon. The picture is much different on the downstream, where Real-Time Entertainment is already a significant portion of daily traffic (23.2%), and is poised for enormous growth in the future. P2P Filesharing represents 21.0% of daily downstream traffic, but this figure will drop as users come to embrace the convenience of on-demand entertainment.

Typical of relatively new broadband markets, Web Browsing is the largest downstream category, and in Africa accounts for 33.5% of daily downstream traffic. Despite this early dominance, as time progresses (and perhaps as early as the Spring 2012 study), we expect to see Real-Time Entertainment growth to outpace that of both P2P Filesharing and Web Browsing, such that it becomes the dominant category by volume.

When we shift the examination to focus on the evening's busiest hours (from 6pm to 11pm in Africa), the upstream traffic profile changes substantially as Real-Time Communication applications surge to become the second-largest consumer of upstream bandwidth. This rise is driven primarily by Skype, which alone represents more than 15% of upstream bytes during the evening.

Monthly consumption shows less disparity than in most mature broadband markets, with median and mean monthly data usage being 2.4 GB and 6.8 GB, respectively. Collectively, the top 1% of heaviest upstream users account for 10% of total upstream usage. By comparison, the lightest 60% of subscribers overall only account for 10% of traffic.

Further Reading

For more information about broadband trends in emerging markets, download Sandvine's Global Internet Phenomena Spotlight: Emerging Markets from sandvine.com

Average Day Network Overview

A first step towards understanding network utilization patterns is to examine the high level ebb and flow of active subscribers and aggregate traffic. Figure 1 illustrates this pattern; both the Active Subscribers and Aggregate Traffic lines are expressed relative to their own daily maximum. To be considered “active”, a broadband connection must have exceeded a minimum usage threshold per hour, while the Aggregate Traffic line is simply the sum of all upstream and downstream subscriber traffic. The shaded portions represent the time period for which the lines are within 5% of the maximum.

Did you know?

The average fixed access network in Africa has a downstream-to-upstream ratio of 5.28

In the African networks examined, the subscriber curve rises rapidly after the 5am minimum to reach a local maximum at noon. It then rises steadily, but more slowly, before attaining its absolute maximum at 4pm. The line is within 5% of this maximum for a total of two hours, which is quite focused by global standards.

The traffic flow lags that of subscriber activity, and peaks at 6pm before hovering near the maximum for almost 5 hours.




Figure 1 - Average Day (Subscribers and Traffic) - Africa, Fixed Access

Figure 2 shows the relative levels of downstream and upstream traffic that contribute to the aggregate traffic curve in Figure 1. Subscriber downloads represent the bulk of the traffic, but it becomes clear that the 6pm traffic maximum observed in Figure 1 is the result of a surge in downstream traffic beginning at 5pm. Interestingly, there does not appear to be a corresponding increase in the upstream component during this surge.




Figure 2 - Average Day (Upstream and Downstream) - Africa, Fixed Access

Figure 3 shows how the ratio of downstream traffic to upstream traffic varies throughout the day. The maximum value of 6.41 is achieved at 2am, which is unusual of a maximum ratio - it is typically achieved in the evening period when particularly asymmetric applications are being used.




Figure 3 - Average Day (Downstream-to-Upstream Ratio) - Africa, Fixed Access

Different categories of traffic have dramatically different downstream-to-upstream ratios. Figure 4 illustrates this reality in detail. The highest ratio (27.11) belongs to Bulk Entertainment, which is typically a unidirectional download (excepting a small amount of control overhead). It is unusual to see Gaming traffic with such a high ratio (25.79). A potential explanation could be that while gaming content downloads remain on par with other regions, perhaps the actual gameplay component (which is more symmetric in nature) is lower, so there is little upstream traffic to bring the ratio down.

The overall ratio of 5.28 reflects a global shift towards a higher proportion of downstream traffic.

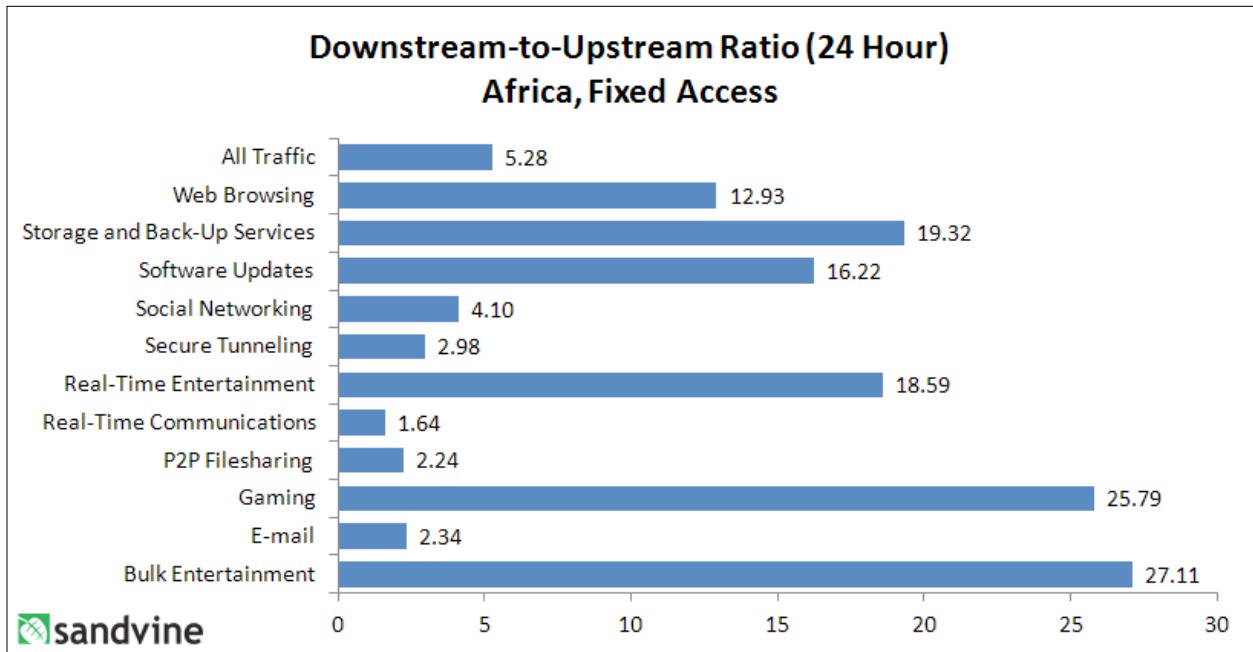


Figure 4 - Downstream-to-Upstream Ratio by Category (24 Hour) - Africa, Fixed Access

With relatively few applications making extensive use of upstream capacity, it is typical that P2P Filesharing is the dominant category; indeed, this characteristic is true for Africa's broadband networks, as shown in Figure 5. P2P Filesharing accounts for 49.4% of daily upstream traffic, with Web Browsing a distant second with 14.5%. There is a notable uptick in P2P activity starting at 5pm, which could be at least partially responsible for the aggregate traffic peak observed in Figure 1, although examination of the downstream patterns is required to confirm this relationship. The daily profile clearly illustrates that Real-Time Communications is the second-largest category during the network's busiest periods, but the drop-off overnight lowers its daily average to 13.7%, good enough for third place. The top 5 categories are rounded out by Real-Time Entertainment (6.6%) and Social Networking (4.7%).

Did you know?

P2P Filesharing accounts for 49.4% of upstream traffic in Africa, far and away more than any other category - this is similar to where the USA was x years ago




Figure 5 - Average Day (Network Upstream) - Africa, Fixed Access

Globally, we have seen over the past few studies a shift in downstream consumption away from traditional Web Browsing and toward Real-Time Entertainment applications - does this trend carry into Africa? Indeed! Figure 6 confirms that there is a substantial amount of Real-Time Entertainment traffic on Africa's broadband networks; 23.2% by daily volume. Nevertheless, Web Browsing remains the leading category, accounting for 33.5% of bytes. P2P Filesharing accounts for 21.0% of downstream traffic in Africa; Figure 6 confirms that the aggregate traffic peak observed in Figure 1 is largely driven by a sudden surge in P2P activity. This peculiar rise was observed across several networks, and is indicative of a widespread phenomenon, rather than being the result of localized network management.

As time progresses, and perhaps as early as the Spring 2012 study, we expect to see Real-Time Entertainment rise and overtake Web Browsing to become the dominant downstream category, primarily at the expense of P2P but also by partially cannibalizing web traffic.

Did you know?

Web Browsing represents 33.5% of downstream bytes, but Real-Time Entertainment is coming on strong with 23.2%

Further Reading

For more insight into the growth of Real-Time Entertainment traffic in Africa, be sure to read "Fast Forward - The future of online video in EMEA", available on sandvine.com.¹




Figure 6 - Average Day (Network Downstream) - Africa, Fixed Access

1. The article originally appeared in Connect-World EMEA, we have reproduced it here: http://www.sandvine.com/downloads/documents/EMEA_2011_03.pdf

Data Consumption Characteristics

While most published Internet consumption statistics address the mean, the median is actually much more representative of the subscriber base as a whole. That is not to say that analysis of the mean is not valuable; rather, it is important to examine not only both the median and the mean, but also the relationship (specifically, the ratio of mean to median) between the two.

Table 1 shows the median upstream, downstream and aggregate consumption per day, week, and month in Africa. The median monthly data consumption is 2.4 GB, but it is not uniformly distributed throughout the entire month, as evidenced by the weekly median consumption of 1.8 GB. These values suggest that most subscribers tend to concentrate their usage. However, daily usage does tend to be uniformly distributed throughout the week, which can be inferred by comparing the weekly consumption metrics to the daily ones.

	Median Consumption - Africa, Fixed Access		
	Day	Week	Month
Upstream	31.0 MB	221.3 MB	299.4 MB
Downstream	250.0 MB	1.5 GB	2.1 GB
Aggregate	295.5 MB	1.8 GB	2.4 GB


SOURCE: SANDVINE NETWORK DEMOGRAPHICS 

Table 1 - Median Consumption - Africa, Fixed Access

Analysis of the mean allows network operators to quantify the requirements of the network's heavier users. Table 2 enables this analysis. Mean monthly consumption is 6.8 GB, and once again the bulk tends to be achieved in a short period of time.

	Mean Consumption - Africa, Fixed Access		
	Day	Week	Month
Upstream	104.9 MB	606.6 MB	1.0 GB
Downstream	614.8 MB	3.4 GB	5.8 GB
Aggregate	719.7 MB	4.4 GB	6.8 GB


SOURCE: SANDVINE NETWORK DEMOGRAPHICS 

Table 2 - Mean Consumption - Africa, Fixed Access

The mean is heavily influenced by the network's highest data consumers, so examining the ratio of the mean to median (Table 3) can indicate if network usage is imbalanced.

The ratios aren't as high as we have observed in past reports for fixed access networks in other regions, suggesting that network usage is more uniformly distributed in Africa than in other regions. Nevertheless, as is consistent with general observations on networks around the globe, the highest ratios appear on the upstream access network (typically, this phenomenon is the result of P2P Filesharing users having relatively high monthly usage).

	Mean : Median Ratio - Africa, Fixed Access		
	Day	Week	Month
Upstream	3.38	2.74	3.47
Downstream	2.46	2.21	2.79
Aggregate	2.44	2.21	2.82


SOURCE: SANDVINE NETWORK DEMOGRAPHICS 

Table 3 - Mean-to-Median Ratio - Africa, Fixed Access

Figure 7 shows the cumulative consumption of subscribers over a month, and illustrates the disparity between the consumption of the network's heaviest users and the majority of subscribers. Overall, the cumulative distribution of data consumption is remarkably consistent between the upstream, downstream, and aggregate in Africa; in fact, the curves overlap almost exactly.

Generally, the top 1% of subscribers in Africa account for about 10% of the total resource utilization, and the top 20% of subscribers account for 65% of traffic. While this imbalance seems high, it is actually lower than what is generally observed elsewhere. On the opposite end of the usage spectrum, the lightest 60% of subscribers generate only 10% of traffic.

By the Numbers

- 2.4 GB** - median monthly aggregate usage
- 6.8 GB** - mean monthly aggregate usage
- 2.82** - monthly aggregate mean-to-median ratio

Did you know?

The top 1% of users in Africa account for 10% of usage, the same amount as the bottom 60%




Figure 7 - Percentile Byte Contribution (Monthly) - Africa, Fixed Access

Focus on Peak Period

Networks must be architected for peak utilization requirements, so a detailed understanding of the factors driving peak period data consumption are of vital importance to communications service providers hoping to optimize their infrastructure investments.

Figure 7 shows the relative contribution of the top 5 traffic categories by bandwidth, to the upstream, downstream, and aggregate totals during the evening period of heaviest traffic. The reader should keep in mind that the total amount of downstream traffic far exceeds the upstream. Therefore this figure illustrates how traffic breaks down within each traffic direction and comparison should be limited to relative quantities, rather than absolutes.

On the upstream, P2P Filesharing is the dominant category, although Real-Time Communications represents almost one quarter of total bytes (much higher than is observed in more mature broadband markets). As previously observed when discussing Figure 5, the Real-Time Communications category surges in the evening, and is not evenly distributed throughout the day.

Did you know?

- Real-Time Communications applications generate 23.2% of Africa's peak upstream demand.
- Despite the relatively early stage of Africa's emergence as a broadband market, Real-Time Entertainment already accounts for more than 20% of peak traffic.

The downstream network has a more balanced ecosystem of categories, and the big three (Web Browsing, Real-Time Entertainment, and P2P Filesharing) together account for 74.8% of traffic volume. Social Networking represents 4.4% of downstream traffic, falling just shy of cracking the top 5 list.

In Africa at least, there appears to be four big categories of interest, consisting of the usual three but with the addition of Real-Time Communications, which accounts for fully 9.5% of peak aggregate demand. The category will likely remain high so long as the telephone infrastructure is immature and international tariffs remain high. Social Networking might well emerge as a partial substitute for voice communications, so we could see it rise beyond the 4.7% of peak aggregate traffic shown in Figure 8. As mentioned previously, expect to see Real-Time Entertainment increase its share of total traffic at the expense of both P2P Filesharing (primarily) and Web Browsing.




Figure 8 - Peak Period Traffic Composition - Africa, Fixed Access

Table 4 shows the Top 10 applications on the upstream, downstream and in aggregate, during peak period. We already have a hint of what will appear on this list, but closer examination often reveals some unexpected facts.

Upstream capacity is consumed primarily by BitTorrent, which is the dominant P2P Filesharing network worldwide. There is a larger amount of eDonkey than is typically observed, but eDonkey is known to be popular with French-speakers, so that could be the explanation for its presence on this list. There is more Skype than all the other voice protocols combined, and Skype is the primary driver of the Real-Time Communications category in Africa.

In the downstream direction, HTTP accounts for almost a third of evening traffic, while BitTorrent generates 17.6% of traffic. On the entertainment side, there is more generic Flash Video than YouTube, which is unusual - expect to see this reverse as more international cables come online to increase Africa's access to global content.² MegaUpload and Megavideo live side-by-side in the list, and symbolize the divide between "watch later" and

Did you know?

- BitTorrent represents more than 20% of peak traffic in Africa, second only to HTTP. Skype is more than 7% of peak, and more than 15% of utilized upstream bandwidth.
- YouTube is only 5% of Africa's peak traffic - expect this figure to rise as new undersea cables come online and increasingly connect Africa to the world's major content sources.

2. Middleton, James. "Cable Landing Boosts African Connectivity." Telecoms.com. Telecoms.com, 11 Oct. 2011. Web. <<http://www.telecoms.com/34593/cable-landing-boosts-african-connectivity/>>.

“watch now” behaviors. As Africa’s networks become more mature and more content becomes available, MegaUpload will fall off this list entirely, to be replaced by Megavideo and other Real-Time Entertainment applications and services.

Facebook is the 6th largest consumer of Africa’s broadband resources in the evening, accounting for 4.42% of aggregate traffic. In general, Africa’s networks are dominated by less sensitive applications, but time will see the increased adoption of real-time services that demand network reliability and consistency.

Rank	Upstream		Downstream		Aggregate	
	Application	Share	Application	Share	Application	Share
1	BitTorrent	38.09%	HTTP	32.20%	HTTP	28.94%
2	Skype	15.33%	BitTorrent	17.58%	BitTorrent	20.81%
3	HTTP	11.50%	Flash Video	6.09%	Skype	7.17%
4	Facebook	5.72%	Skype	5.64%	Flash Video	5.26%
5	eDonkey	4.54%	YouTube	5.56%	YouTube	4.96%
6	RTP	4.30%	Facebook	4.18%	Facebook	4.42%
7	Teredo	2.13%	MegaUpload	3.53%	MegaUpload	3.11%
8	SIP	2.02%	Megavideo	2.69%	Megavideo	2.33%
9	SSL	2.00%	iTunes	2.49%	iTunes	2.16%
10	YouTube	1.77%	RTMP	2.09%	RTMP	1.83%
	Top 10	87.40%	Top 10	82.05%	Top 10	81.00%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS




Table 4 - Top Peak Period Applications by Bytes - Africa, Fixed Access

The evening differs from the rest of the day in significant main ways:

- More subscribers are online, so traffic levels increase.
- Subscribers who are online are more active; that is, they are using more bandwidth individually than at any other time of day.
- Subscribers who are online tend to favor on-demand and interactive applications, the very applications most sensitive to network issues.

The result is a perfect storm of more people using more sensitive activities at the time when the network is most heavily taxed. Consequently, network operators closely monitor metrics that correlate with the quality of the subscriber experience.

One such metric is the Prime Time Ratio, of which two versions are popular:

1. Network Prime Time Ratio: compares the absolute levels of network traffic during an average peak period hour with an average off-peak hour. This metric is most helpful for capacity planning
2. Subscriber Prime Time Ratio: compares the relative level of individual subscriber traffic during peak period with the off-peak period. This metric is most helpful as an indicator of subscriber quality of experience (a ratio higher than 1 means that each active subscriber is able to use more bandwidth during the evening)

The Network Prime Time Ratio for all traffic is 1.38, so the average peak period hour is 38% busier than the average off-peak hour. In positive news for Africa’s subscribers, the network is still able to deliver an overall Subscriber Prime Time Ratio of 1.11, despite the increased load.

The biggest causes of the evening surge are Real-Time Communications and Social Networking, since they have high ratios coupled with relatively high volumes. By contrast, while Bulk Entertainment and Gaming traffic have high ratios, they don’t represent as much absolute traffic. Nevertheless, it is clear that Africa’s subscribers enjoy an evening gaming session.

By the Numbers

1.38 - Network Prime Time Ratio

1.11 - Subscriber Prime Time Ratio

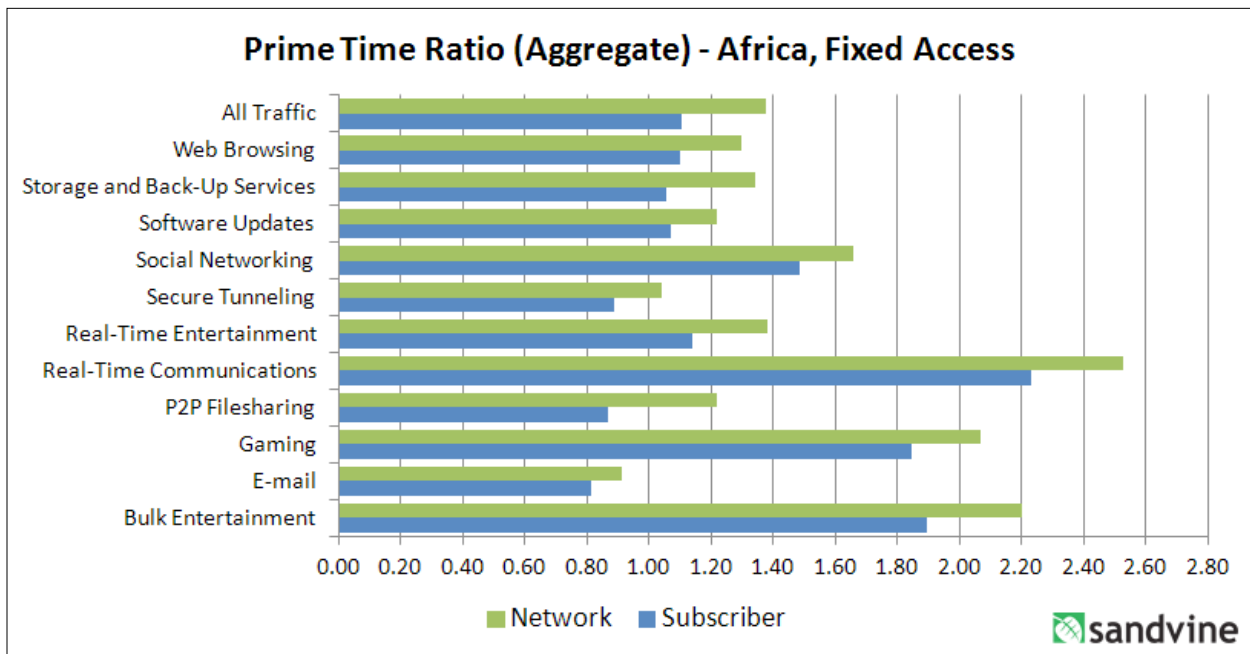


Figure 9 - Prime Time Ratio (Aggregate Traffic) - Africa, Fixed Access

Study Details

Sandvine's Global Internet Phenomena Reports examine a representative cross-section of the world's leading fixed and mobile communications service providers and are made possible by the voluntary participation of our customers. Collectively, Sandvine's customers provide Internet and data service to hundreds of millions of subscribers worldwide.

The data gathered for these reports is completely subscriber-anonymous. No information regarding specific content or personally-identifiable information (including, but not limited to, IP or MAC addresses and subscriber IDs) was collected during this study.

The overall study examined Internet traffic in September 2011 from five regions:

- North America
- Latin America
- Europe
- Asia-Pacific
- Africa

This study reflects the traffic profiles of real service providers, including the impact of any network management (for instance, congestion management and traffic optimization) policies that may be in place.

The data collected includes the bandwidth per second per protocol and the number of active hosts per protocol on the network at each hour. Data also includes the total transmitted (upstream) and received (downstream) bytes, from the subscriber's perspective, attributable to each subscriber for the 30 days, 7 days, and 1 day preceding the time of data collection.

The datasets were used to create a 24-hour profile of each network, normalized by the number of active subscribers at each hour in the day. These profiles were then aggregated hierarchically for each region with weightings based on subscriber counts and access technology market share.

The transmitted and received bytes per subscriber data sets were used to create ordinal rankings of all subscribers on a network based on a combination of data direction (upstream, downstream, aggregate) and data period (day, week, month), for a total of nine ranked lists ordered by total byte usage. These lists enabled consumption analysis based on percentile ranking and cast light on the widely varying data needs of individual subscribers.

In parts of the report we reference industry publications, analyst studies, media articles and other sources. As such, we are indebted to the collective work and wisdom of a large number of individuals and organizations and have endeavoured to correctly cite all sources and to identify the original creator of referenced material.

Explanation of Categories

The table below describes each of the categories used in the *Global Internet Phenomena Spotlight - Africa, Fixed Access, Fall 2011*.

Category	Description	Example Applications and Protocols
Anonymity	Protocols that mask or obfuscate application or individual identity	Tor (The Onion Router)
Bulk Entertainment	Entertainment that is acquired in bulk then consumed sometime after arrival	iTunes, movie download services
Bulk Transfers	Large data transfers using the File Transfer Protocol or its derivatives	FTP (File Transfer Protocol)
E-mail	Service-provider and webmail e-mail services	SMTP, POP3, webmail (Hotmail, Gmail, etc), BlackBerry encrypted e-mail
Encapsulation Tunnelling	Tunnels used for wrapping traffic	L2TP, GRE, Teredo, 6 to 4
Gaming	Console and PC gaming, console download traffic, game updates	Nintendo Wii, Xbox Live, Playstation 2, Playstation 3, PC games (for example, World of Warcraft)
Network Administration	Protocols and services used to administer the network	DNS, ICMP, NTP, SNMP
News Groups	Network news services (where “news” means “data” - it doesn’t have to be actual news)	NNTP, encrypted NNTP (over SSL)
P2P Filesharing	File-sharing applications that use a peer-to-peer distribution model	BitTorrent, eDonkey, Gnutella, Ares, Winny, Share, Foxy, Pando
Real-Time Communications	Applications and protocols that allow interactive chat, voice, and video communications	Skype, MSN Messenger, ICQ, SIP, MGCP, AOL Instant Messenger (AIM), IRC, Oovoo, Jabber, Gadu-Gadu, MGCP, Facetime
Real-Time Entertainment	Applications and protocols that allow “on-demand” entertainment that is consumed (viewed or heard) as it arrives	Streamed or buffered audio and video (RTSP, RTP, RTMP, Flash), peercasting (PPStream, Octoshape), placeshifting (Slingbox, home media servers), specific streaming sites and services (Netflix, NCAA, Hulu, YouTube, Google Video, Spotify, BBC iPlayer)
Remote Connectivity	Protocols and services that allow remote access to network resources	Remote Desktop, VNC, PC Anywhere
Secure Tunnelling	Encrypted tunnels typically used for Virtual Private Networks and secure web transactions	SSL, SSH
Social Networking	Websites and services focused on enabling interaction (chat, communication) and information sharing (photos, status, etc) between users	Facebook, MySpace, Twitter, Habbo, Bebo
Software Updates	Application updates for software, firmware, and operating systems	Windows Update, anti-virus updates
Storage and Back-Up Services	Services that provide file-hosting, network back-up, and one-click downloads	PDBox, Netfolder, Rapidshare, MegaUpload, Mozy, zShare, Carbonite, Dropbox
Web Browsing	Web protocols and specific websites	HTTP, WAP browsing

Bibliography

Middleton, James. "Cable Landing Boosts African Connectivity." Telecoms.com. Telecoms.com, 11 Oct. 2011. Web. <<http://www.telecoms.com/34593/cable-landing-boosts-african-connectivity/>>.

Headquarters
Sandvine Incorporated ULC
Waterloo, Ontario Canada
Phone: +1 519 880 2600
Email: sales@sandvine.com

European Offices
Sandvine Limited
Basingstoke, UK
Phone: +44 (0) 1256 698021
Email: sales@sandvine.co.uk

Copyright ©2011 Sandvine
Incorporated ULC. Sandvine and
the Sandvine logo are registered
trademarks of Sandvine Incorporated
ULC. All rights reserved.



sandvine[®]
Intelligent Broadband Networks