XENA NETWORKS

WHITE PAPER

MICROBURST

OVERVIEW

Microburst is a phenomenon where data packets are transmitted in rapid burst. It can lead to periods of full line-rate transmission that overflow packet buffers of the network stack. Traffic flows and applications often show bursty behavior when transported across IP networks or in data centers. Microbursts, though short-lived, are difficult to detect and monitor and can cause network and application performance degradation, i.e. increased latency, jitter, and packet loss.

Given that microbursts are difficult to detect and have great negative impact on network and application performance and can result in business loss, enterprises networks, data centers, and network elements should test their ability to handle microbursts on different layers.

With its Valkyrie and Vulcan test platforms, Xena provides advanced microburst test solutions from stateless Layer 2-3 to stateful Layer 4-7 with different burst profiles at different link speeds, bandwidth utilizations, etc. Different burst profiles, i.e. burst size, packet size, and burst density, have different impacts and should therefore all be tested.

"Microbursts can seriously impact network performance but are difficult to detect. Xena provides advanced microburst test solutions for Layer 2-7."

XENA NETWORKS

WHITE PAPER

MICROBURST TESTING

Contents

Introduction

What is microburst?

Mircoburst is a phenomenon where rapid bursts of data packets are transmitted in quick succession. This can lead to periods of full line-rate transmission that can overflow packet buffers of the network stack. When measuring traffic data rate with a small interval size, e.g. 1 millisecond, microbursts become visible. Figure 1 shows that when reduce the interval size to 1 ms, microburst become obvious while it is hard to detect with a larger interval size.

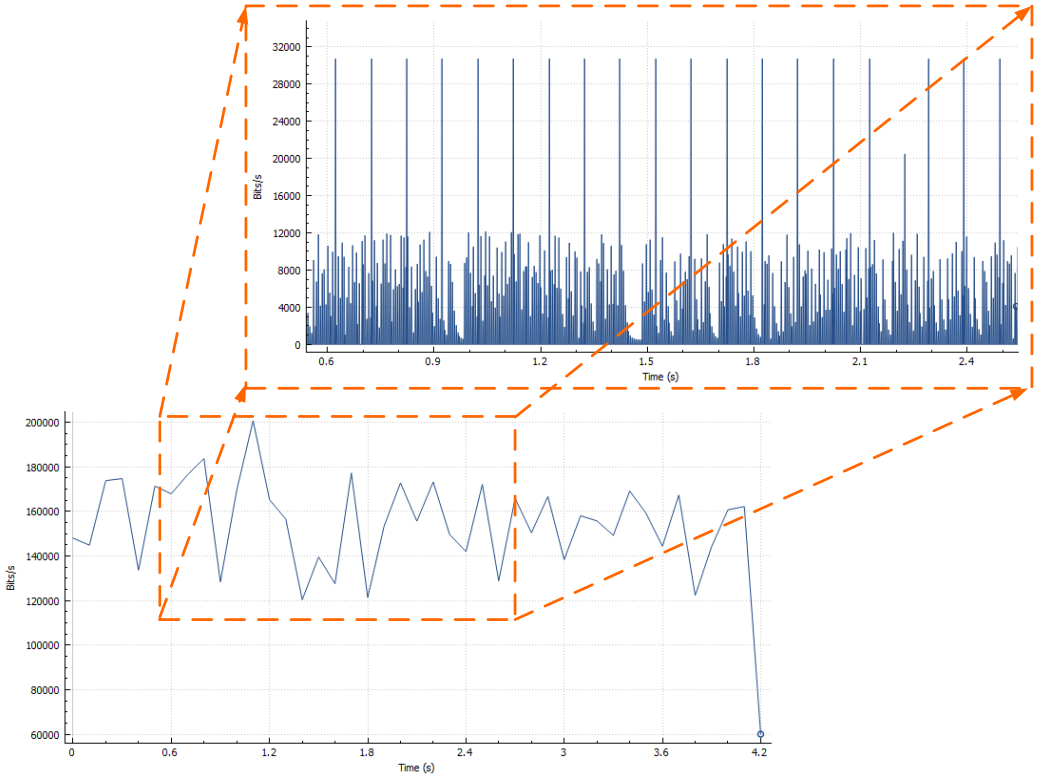

Figure 1: Periodic microbursts

What causes microbursts?

Traffic flows and applications often show bursty behavior when transported across IP networks or in data centers. Burstiness is due to the nature of packet-oriented communication, including packetization and packet handling processes in switches and routers. It can also originate from some applications that are designed to provide burstiness to the networks, e.g. data pushing, IPTV, etc. Since those bursts occur over short periods, they are usually referred to as microbursts.

IPTV application is such an application that introduce microburst as designed. Videos are distributed to users in frames. Frames are periodic around the frame rate, the frequency

(measured in frames per second, fps) at which an imaging device displays consecutive images. For example at 30 fps, a frame is sent very 33 ms. The size of a frame can vary greatly due to video compression techniques, where I-frames, P-frames, and B-frames are used to compress data and save bandwidth. I-frames are full reference frames and they are larger than the P-frames. When I-frames are sent, they look like a small burst. Since frames are sent in groups of packets, an I-frame may take as many as 80 packets or more for a high-definition 1080 video. These 80+ packets will create a microburst on the link.

IPTV application is such an application that introduce microburst as designed. Videos are distributed to users in frames. Frames are periodic around the the frame rate, the frequency (measured in frames per second, fps) at which an imaging device displays consecutive images. For example at 30 fps, a frame is sent very 33 ms. The size of a frame can vary greatly due to video compression techniques, where I-frames, P-frames, and B-frames are used to compress data and save bandwidth. I-frames are full reference frames and they are larger than the P-frames. When I-frames are sent, they look like a small burst. Since frames are sent in groups of packets, an I-frame may take as many as 80 packets or more for a high-definition 1080 video. These 80+ packets will create a microburst on the link.

Figure 2: IPTV application and microburst

Another example is TCP incast. TCP incast is a many-to-one (or sometimes referred to as fan-in) communication commonly found in data centers for cloud applications. For example in Figure 3, a single parent server places a request for data to a group of nodes that receive the request simultaneously. The group of nodes respond to the parent server in a synchronous fashion and result in a microburst of a many-to-one TCP data stream scenario.

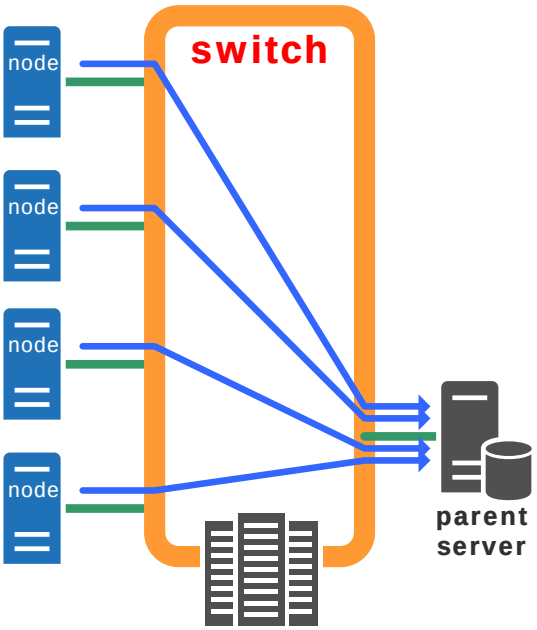

Figure 3: TCP incast and microburst

WHITE PAPER

Impact of microbursts

Microbursts, though short-lived, are difficult to detect and monitor and can cause network and application performance degradation, i.e. increased latency, jitter, and packet loss. In a fan-in communication scenario where many send data to one, packets are stored in the memory and latency will obviously be increased. This kind of latency increase (network latency) is due to the longer waiting time of each packet in the queue caused by congestion and can happen to any switch and router in the network.

Another kind of latency (application latency) caused by packet loss is more serious. As shown in the example in Figure 3, the simultaneous many-to-one burst can cause egress congestion at the network port attached to the parent server. If the switch cannot handle microbursts well, packet loss will occur. The nodes will detect the loss by ACK timeout, retransmit the data, and ramp up the throughput according to the TCP congestion control procedure. Thus, the application will have to wait until all the data is received or simply complete the job with partial data fetched. Either way is a great degradation to the application performance. For financial trading transactions, such a sudden performance degradation caused by microburst could result in a great investment loss.

Microbursts can also increase traffic jitter because the impact can be forwarded through various network elements. Switches with slow clock rates can show very high jitter with a microburst that suddenly increases the queue length.

Another impact is simultaneous arrival of microbursts, as shown in Figure 4. Simultaneous arrival means that microbursts appear on multiple ports of a switch/router at the same time. Usually the switch fabric has the bandwidth capacity to process traffic from all the ports. However, when sudden microbursts occur on multiple ports, they put the internal buffers to pressure and may cause internal buffer congestion and further lead to out-of-sequence, latency increase, or even frame loss.

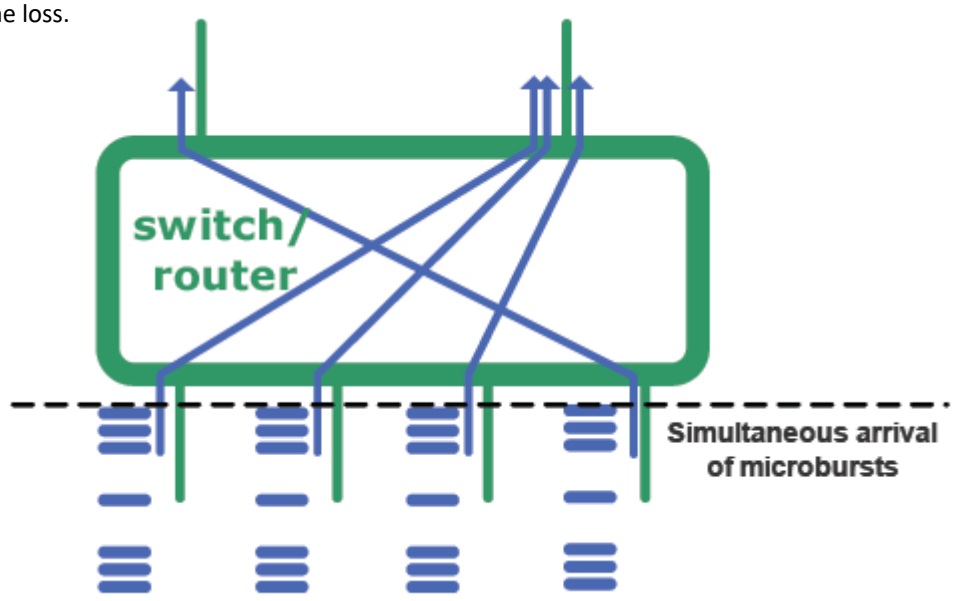

Figure 4: Microbursts simultaneous arrival

Test Microburst with Xena

Given that microbursts are difficult to detect and have great negative impact on network and application performance and can result in business loss, enterprises networks, data centers, and network elements should test their ability to handle microbursts on different layers. For example, buffer statistics are extremely important for congestion analysis in SDN/NFV and other scenarios on both network and application layers. In order to verify that the trigger-based mechanism can detect microbursts, a rich test solution is required on different layers. Devices such as packet brokers are developed to mitigate microburst with high data burst buffers, and should be thoroughly verified. Xena is capable of providing advanced microburst test solutions from Layer 2 to Layer 7 with different burst profiles at different traffic rates.

Test Microburst on Layer 2-3

Different burst profiles, i.e. burst size, packet size, and burst density, could have different impacts on Layer 2-3 QoS/performance and therefore should all be tested. Figure 5 shows different burst profiles.

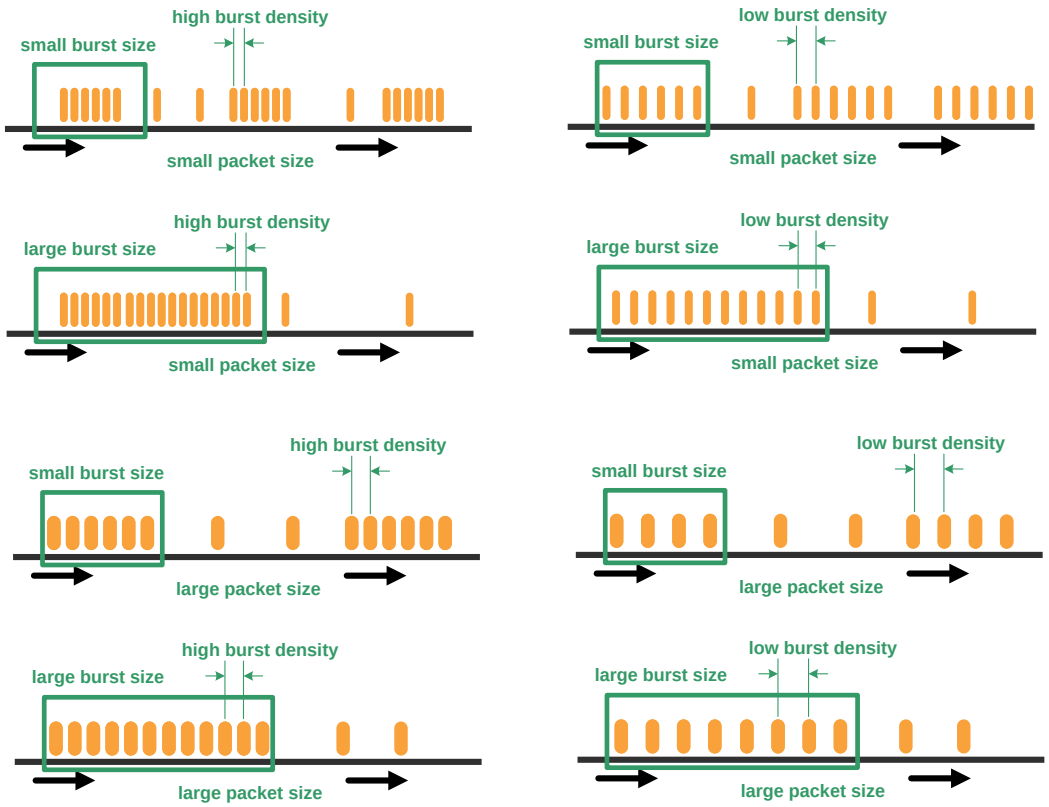

Figure 5: Different burst profiles for microburst testing on layer 2-3

Microbursts usually only last a short period of time and hard to detect. Thus, it is import to test whether the device/system under test (DUT/SUT) is capable of detect/monitor microbursts. A sudden burst in traffic flows may cause buffer congestions and reduce the QoS performance, such as latency, jitter, and frame loss. Thus, all the flows must be monitored. In addition, a conforming traffic flow can be affected by a flow that generates microbursts. This cross-flow effect should also be examined since some switches and routers do not provide flow isolation. Figure 6 shows an example of Xena generating microbursts with a conforming traffic flow.

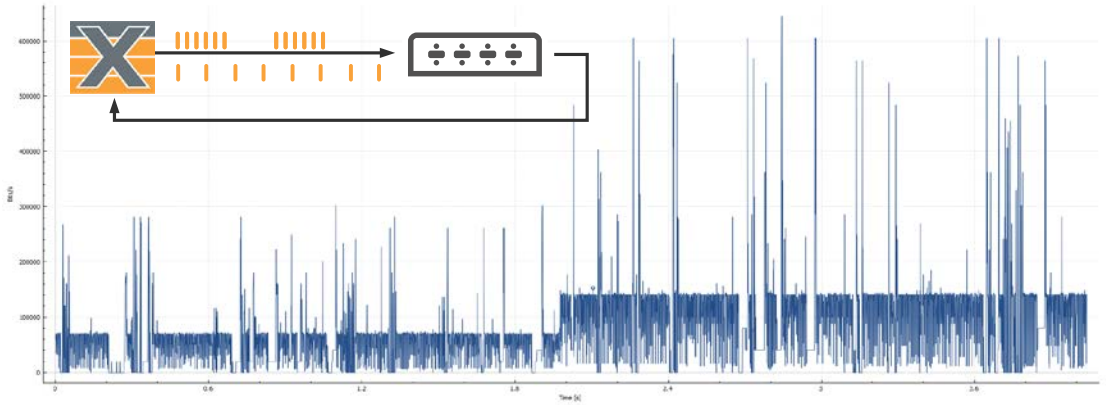

Figure 6: Xena microburst testing on stateless Layer 2-3, 1ms interval

Test microburst on Layer 4-7

Creating microburst traffic not only on stateless Layer 2-3 but also on stateful Layer 4-7 with Xena provides an all-round validation of DUT/SUT's microbursts handling ability. Within a TCP connection, bursts can be created periodically with controlled duration and frequency of the bursts. As shown in Figure 7, users can define burst percentage by configure burst duration and inactive duration. In addition, different TCP segment sizes can be applied to the test.

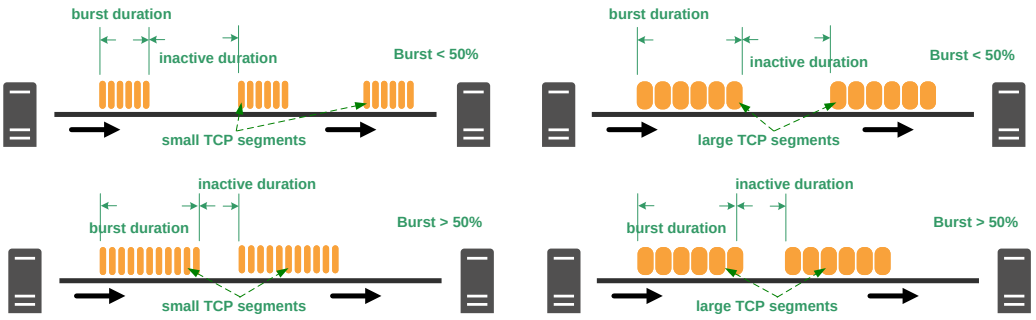

Figure 7: Different burst profiles for microburst testing on layer 4-7

Packet loss can heavily affect the performance of TCP because the sender will have to retransmit and ramp up the throughput according to the TCP congestion control procedure Thus, it is vital to verify that the DUT/SUT can process and forward 100% of the traffic with zero loss under

WHITE PAPER

different Layer 4-7 microburst profiles. Figure 8 shows an example of TCP microbursts generated by Xena.

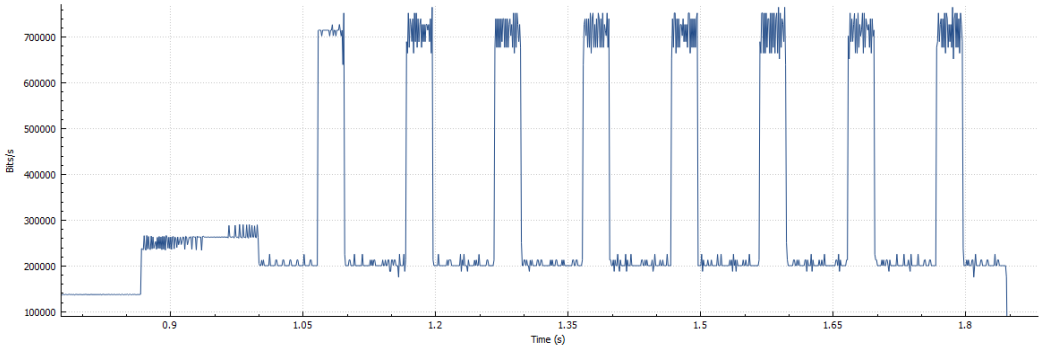

Figure 8: Xena microburst testing on stateful Layer 4-7, 1ms interval.

Simultaneous microbursts arrival test

Microburst simultaneous arrival should also be tested to examine the burst handling capability of the DUT. Xena is able to provide synchronous traffic across multiple ports of the DUT, as shown in Figure 9, which makes sure that microbursts arrive on different ports at the same time. In the example shown in Figure 9, with the charting function shows that how the QoS of conforming streams are affected by the microburst that arrive simultaneously.

WHITE PAPER

Microbursts over WAN tests

With Xena L4-7 testers, impact of microburst traffic to WAN scenario can be tested. As shown in Figure 10, two financial institutions are connected to a Xena L4-7 tester. Stateful TCP connections are established between the institutions across different firewalls, and traffic is transmitted from one end to the other. Together with the normal traffic, microbursts can be injected to test the microburst handling ability of the system. TCP retransmissions caused by packet loss will monitored and reported. Since TCP retransmissions result in latency increase and QoS degradation, automated iterative tests should be carried out in order to generate a clear picture of the capability of the system under test.

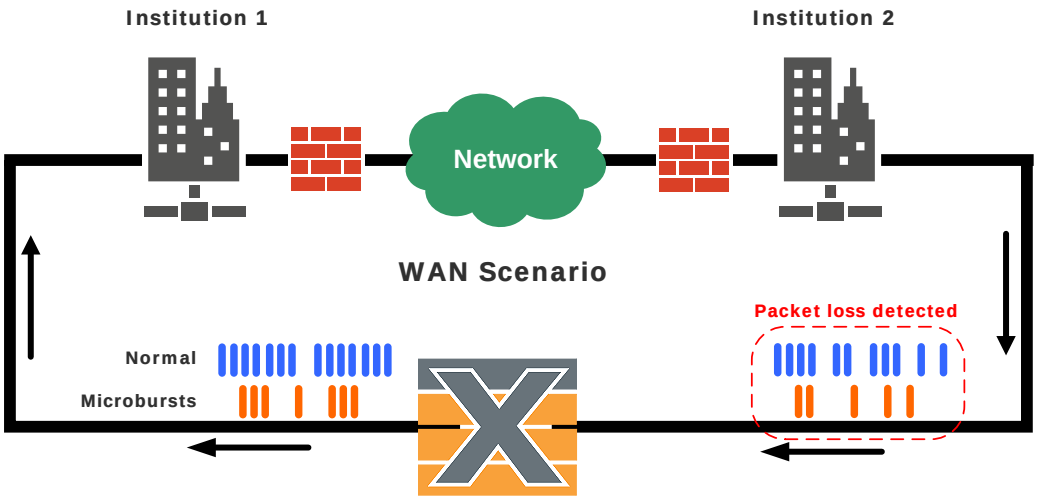

Figure 10: Test impact of microburst over WAN.

CONCLUSION

Microbursts traffic can cause packet drops that can further bring serious performance degradation, e.g. TCP incast problem, increase latency and jitter. Since microbursts usually last only a short period, they are difficult to detect and monitor. If the DUT/SUT is not tested for its ability to handle microbursts, great business loss and experience degradation can occur after deploying such systems into the production environment. It is vital to validate microburst-handling ability not only on Layer 2-3 but also on Layer 4-7, in order to have an all-around examination. With its Valkyrie and Vulcan test platforms, Xena provides advanced microburst test solutions from stateless Layer 2-3 to stateful Layer 4-7 with rich burst profiles at different link speeds, bandwidth utilizations, etc. Different burst profiles, i.e. burst size, packet size, and burst density, have different impacts and therefore should all be tested.