



Corvil Insight

Low-Latency
Market Data

Contents

How much bandwidth do you really need?	4
How to ask the bandwidth question for low-latency?	4
What about loss?	5
Exchange feed example	5
Working with exchange and data provider recommendations	7
Looking at the document in more detail	9
Before you act – check the following	10
Further Reading	10

How much bandwidth do you really need?

Achieving low-latency distribution of market data has become one of the most important factors that determine success in high frequency trading. No matter how fast you can execute, if your market data is delayed relative to competing traders, you will not achieve the expected fill rates.

The key to achieving low-latency for market data distribution is to ensure sufficient bandwidth has been provisioned in the access network to the market centers of interest. But many questions arise:

- Is a 10Gbit/s connection or co-location worth it?
- How much better will it be over a 1Gbit/s direct connection?
- If I can only get a 100Mbit/s extranet connection how much latency am I going to experience?

These are fundamental and important questions for any high frequency trading organization. Unfortunately, it has been our experience that very few exchanges, market data providers, and telecom service providers can answer these questions in a quantitative manner. Typical answers come in the form:

- We recommend a 10Gbit/s LAN connection in our co-location center because it offers the lowest latency possible

If your market data is delayed relative to competing traders, you will not achieve the expected fill rates.

- We have measured the latency of our 1Gb/s access links and have found them to be less than 1 millisecond

The problem with these answers is that they are not useful. They have a “feel good” factor but not the basis for making a rational decision on the purchase of bandwidth for market data.

How to ask the bandwidth question for low-latency?

Our starting point is to make sure we are asking the bandwidth question in the right way. We must first realize we are looking for the bandwidth, which gives us the desired latency performance. Therefore asking the question of how much bandwidth is needed without specifying the latency objective is inviting a vague answer. Second, the bandwidth needed to deliver a specific latency objective is directly dependent on traffic load. If the traffic load spikes, then latency will increase. Therefore the bandwidth required must deliver the specified latency objective even in the presence of traffic spikes, i.e. microbursts. Last but not least, to fully express the question, you need to specify the compliance level of the requirement? For example, do you want the bandwidth to deliver the latency objective for 100% of packets or perhaps 99.99% of packets is sufficient. Our experience shows that this 0.01% delta in compliance factors can result in significantly different bandwidth requirements and hence a lot of money can be saved.

Therefore, an unambiguous and fully specified way to ask the bandwidth question for a market data feeds is:

“For this feed, how much bandwidth is needed for no more than 100 microseconds of latency to be experienced by 99.99% of packets during the busy 1 second of the trading day”

In this example:

- 100 microseconds is the latency objective
- 99.99% is the compliance level i.e. confidence level
- 1 second is the busy period for the service level objective, i.e. corresponds to the maximum message rate or packet rate
- Trading day is the period of time over which the service level agreement applies. This could also be the trading week.

In Corvil we refer to this method of describing bandwidth requirement for low latency as the Corvil Bandwidth (CB) where the latency objective, compliance level, busy period and measurement interval are input parameters that can be specified by the end user.

What About Loss?

When distributing market data, it is equally important to protect against packet loss. Unfortunately, it is not unusual to see network service providers utilize small buffers in routers and switches in an attempt to minimize queuing latencies. This comes at the expense of increased packet loss when microbursts occur. Lost packets cause gaps in market data feeds and ultimately results in much larger latencies being experienced in line handler processing.

Therefore, when asking the bandwidth question it is important to specify a loss objective as well as a latency objective if you want to optimize market data performance. For example:

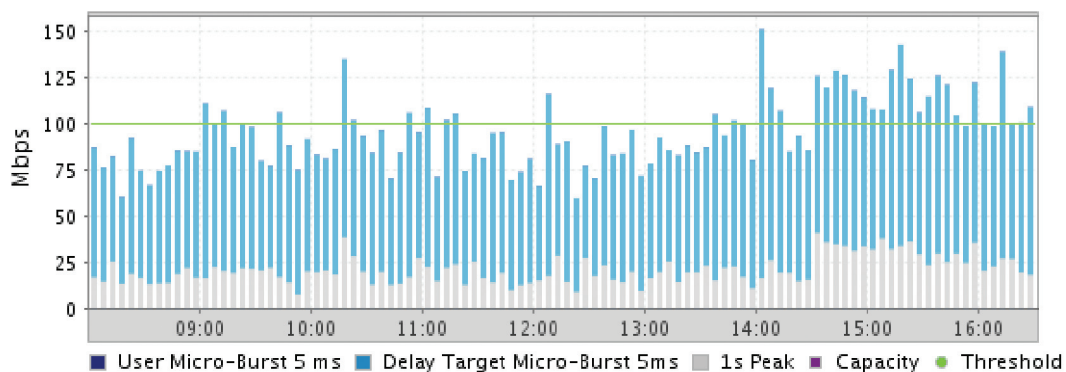
"For this feed, how much bandwidth is needed to protect 99.99% of packets from loss with no more than 100 microseconds of latency to be experienced during the busy 1 second of the trading day"

We recommend using this type of structure when looking to specify the bandwidth requirement and service level agreements with data and network connectivity providers. The Corvil Bandwidth technology referenced above and used in our CorvilNet platform computes the Bandwidth Requirement (BWR) for both latency and loss targets. Two bandwidth numbers are produced, one for latency and one for loss. The highest of the numbers should be used when provisioning a circuit. When the Latency BWR dominates then it is useful to add bandwidth, but when the Loss BWR number dominates, it is often more cost effective to increase buffer sizes rather than adding bandwidth.

Exchange Feed Example

To illustrate the challenges of sizing bandwidth appropriately for low-latency market data feeds we have recorded and analyzed many such feeds at source. An example traffic trace from a leading exchange is shown in Figure 1.

Figure 1
Market data traffic trace over trading day from leading exchange showing 1 sec and 5 millisecond traffic rates






Figure 2
Market data traffic trace for the busy 1 second of the trading day using 1-millisecond microburst traffic rates.

This traffic was destined for a 100Mb/s client access connection.

The first thing to observe from Figure 1 is the significant difference in magnitude between the 1-second rates and 5-millisecond rates. The maximum 1-second rate is 27Mb/s, whereas the maximum 5millisecond rate is 150Mb/s. In fact, if we reduce the granularity to 1-millisecond the microburst traffic rate continues to increase as shown in Figure 2.


The maximum microburst traffic rate during the busy 1-second is now 188Mb/s. As we continue to decrease the measurement granularity the microburst traffic rates will continue to increase. Microburst traffic measurement is a useful indicator of the likelihood of queuing latency happening at congestion points in the network. In this case, it is clear that queuing latency and possible loss will be induced when this traffic hits the 100Mb/s client access connection.

As we continue to decrease the measurement granularity the microburst traffic rates will continue to increase

What is not evident from Figure 2 is how much queuing latency will be induced. Nor is there any information on how much bandwidth is required to assure a specific latency and loss objective can be met. It was this observation, which motivated us to invent the Corvil Expected Quality and Corvil Bandwidth technologies.

With Corvil Expected Quality, we can look at any traffic profile and predict the level of queuing latency and loss that will be induced when it hits a connection of a certain bandwidth. The Expected Latency graph for the market feed on a 100Mb/s connection is shown in Figure 3.

Figure 3
Corvil Expected Latency of the market data feed for the busy 1 second of the trading day on a 100Mb/s client access link.



Latency Target	100 us	1 ms	10 ms
Latency BWR (busy 1 sec)	499 Mbps	132.3 Mbps	54.1 Mbps

From this graph we get a more complete picture. We now see that the 188Mb/s microburst will induce 1.2ms of queuing latency on a 100Mb/s circuit.

Corvil Expected Quality and Corvil Bandwidth can be used to develop bandwidth requirement and expected quality tables for market data feeds showing the sensitivity between bandwidth and latency or loss. This is ultimately what we are looking for when trying to decide on the level of bandwidth to provision. This is shown in Table 1 and Table 2 respectively.

Table 1 tells us if we want 99.9% of packets to experience no more than 100 microseconds of queuing latency on the access network, then 499Mb/s of bandwidth needs to be provisioned.

Corvil expected quality allows us to explore the sensitivity of latency to increasing bandwidth as shown in Table 2.

Table 2 tells us if we upgrade from a 100Mb/s connection to a 1Gb/s connection then the maximum latency over 99.9% of packets will be reduced from 1.206ms to 12us. This allows us to quantify the latency improvement when moving to higher bandwidth connections.

Table 2
Expected Quality table for queuing latency over the busy 1 second of the trading day for 99.9% of packets using Corvil Expected Quality

Link Bandwidth	50 Mbps	100 Mbps	1 Gbps
Latency (Avg)	461 us	139 us	2 us
Latency (Max)	12.84 ms	1.206 ms	12 us

Table 1
Latency BWR table during the busy 1 second of the trading day for 99.9% of packets using Corvil Bandwidth

This allows us to quantify the latency improvement when moving to higher bandwidth connections

Working with exchange and data provider recommendations

During the course of our work with high frequency market participants we have had the opportunity to review and work with the connectivity and bandwidth recommendations from many of the leading exchanges and providers of market data. One of the better descriptions we have found is provided by BATS in their "BATS US Equity/Options Connectivity Manual (Version 5.0.2)". The document emphasizes appropriately some of the challenges involved and provides useful advisory data to the consumer of the market data feeds.

However, questions may still remain for high frequency trader and opportunities exist for further improvement of the type of information presented.

BATS emphasizes firstly that bandwidth recommendations provided are minimum levels (with appropriate emphasis on minimum)

Figure 4 provides an excerpt from the manual outlining the recommendations for bandwidth against specific feeds offered. BATS emphasizes firstly that bandwidth recommendations provided are minimum levels (with appropriate emphasis on *minimum*). Corresponding serialization delays (which we assume includes queuing delays) for the Gigabit and WAN shaped multicast PITCH feed are shown in Figure 5.

The table provided in Figure 4 is useful from an illustrative perspective as it shows both the bit rate and message rate at different measurement intervals for the feeds.

Figure 4

Example illustration of bandwidth recommendations for market data feeds used in high frequency trading. Excerpt is from "BATS US Equity/Options Connectivity Manual (Version 5.0.2)"

3 Bandwidth

3.1 Market Data Feeds

BATS offers four different types of market data feeds:

- Multicast PITCH
- TCP PITCH
- TCP FAST PITCH
- TOP

BATS requires that members allocate a *minimum* of 1 Gb/s per Multicast PITCH Gig-Shaped feed and 100 Mb/s per Multicast PITCH WAN-Shaped feed. With respect to TCP PITCH, FAST PITCH and TOP feeds, BATS understands that firms will have varying levels of sensitivity with respect to latency and as such encourages members to use the statistics provided below to make a well informed decision regarding the bandwidth they will require based on their organization's latency sensitivity.

The table below shows the bandwidth statistics for historical highs for all BATS market data feeds. The table shows the bandwidth peaks for 1, 5, 10, 30, and 60-second intervals. Multicast PITCH peak statistics include 1 and 10 millisecond interval peaks. The TCP statistics include 11 bytes for TCP/IP overhead per packet and do not include the data link layer overhead.

Interval Seconds	Multicast PITCH		TOP		TCP PITCH		TCP FAST PITCH	
	MPS	Mb/s	MPS	Mb/s	MPS	Mb/s	MPS	Mb/s
.001	3,408,000	382						
.010	2,641,900	296						
1	250,861	86	116,063	31	288,612	134	288,612	57
5	132,879	46	63,030	16	162,909	75	162,909	34
10	109,164	38	53,565	14	139,687	65	139,687	33
30	71,136	25	42,773	11	115,783	54	115,783	27
60	60,189	21	36,875	10	100,415	47	100,415	23

*Statistics as of 12/23/09.

It should be noted that BATS data will have microbursts within the one-second interval above, and that these microbursts will exceed the peak rates at the one-second interval. This is demonstrated within the 1 and 10 millisecond interval Multicast PITCH statistics. The extent to which the network connection to the member will cope with the microbursts exceeding the available bandwidth without packet loss will depend heavily on the buffers in the end to end

During spikes in quote updates, members using less than sufficient bandwidth will experience queuing of their market data. Members using the same bandwidth to both receive quotes and transmit orders may expect their orders to be slightly delayed if they have less than sufficient bandwidth. Many members will find these delays unacceptable and should provision their bandwidth to reduce these delays. The following table demonstrates statistics regarding latency incurred as a result of queuing on Gig-Shaped and WAN-Shaped Multicast PITCH feeds.

	Gig-Shaped Multicast PITCH Serialization Delay (ms)	WAN-Shaped Multicast PITCH Serialization Delay (ms)
Average	0	.022
Standard Deviation	.001	1.105
Historical High	5.6	552

* Statistics as of 12/23/09.

**6 of the 12 Gig-Shaped units have never incurred latency due to serialization to date*.

As the volume on the BATS Exchange increases, the market data feed bandwidth required to accommodate peaks will also grow. Members can obtain the latest published market data bandwidth and serialization statistics within this Connectivity Manual. Additionally, monthly statistical updates are presented through the FIF Market Data Capacity working group.

Figure 5

Example illustration of serialization delay measurements for market data feeds used in high frequency trading. Excerpt is from "BATS US Equity/Options Connectivity Manual (Version 5.0.2)"

As expected, the rates increase as the measurement interval reduces. 1 millisecond and 10 millisecond rates are provided for the multicast PITCH feed. As expected there is a significant increase in the peak rates. The one-millisecond peak rate is 382 Mb/s compared to the one-second rate of 86 Mb/s. This gives us a rough estimate of bandwidth needed to achieve latencies in the 1ms ballpark. Predictive analysis technologies like Bandwidth Requirement and Expected Quality can be used to give a more accurate assessment of required bandwidth, expected latency and expected loss.

Looking at the document in more detail

"It should be noted that BATS data will have microbursts within the one-second interval above, and that these microbursts will exceed the peak rates at the one-second interval. This is demonstrated within the 1 and 10 millisecond interval Multicast PITCH statistics.

Our solution also incorporates a broad range of advanced analytics that both parties can use to assess network or application behaviour

The extent to which the network connection to the member will cope with the microbursts exceeding the available bandwidth without packet loss will depend heavily on the buffers in the end to end path."

The challenge highlighted for the end-customer is how to deal with microbursts while assuring latency and loss service levels. The publication of a BWR table for latency and loss (with various buffer sizes) would be most helpful in this situation. Alternatively, market participants could measure and characterize the received feeds themselves using the techniques outlined in the previous sections of this document.

The publication of historical serialization delay information in Figure 5 is unfortunately not very common, but is very useful. It highlights that the maximum latency can be up to 5.6ms for the Gig-Shaped feed and up to 552ms for the WAN-Shaped feed. While these numbers

may look large, they must be interpreted in context, i.e. there is a low probability of them occurring. What would be most helpful is a quantification of this probability (or risk) over the critical trading period, e.g. 99%, 99.9% and 99.99% for the trading day or the busy 1 sec of the trading day. In addition, a predictive "what-if" Expected Latency and Expected Loss table for various link bandwidths would provide useful additional insight and greater confidence when consuming these feeds.

It should be emphasized that the BATS recommendations and discussion is illustrative of the "state of the art" of what we have seen from exchanges on this topic. The recent additions of millisecond microburst and serialization delay measurements contribute important information for the end consumer. As the high frequency markets mature, we expect these type of measurements plus predictive analysis for latency, loss and bandwidth becoming standard requirements.

Excerpts taken from BATS US Equity/Options Connectivity Manual (Version 5.0.2)

BATS Exchange, Inc. ©2009

Before you act – check the following

When signing up for bandwidth connections to market centers for low-latency market data we recommend the following:

1. Make sure you ask the bandwidth question in the right way?

You should specify the latency and loss objectives, the confidence level, the busy period and the period over which the numbers apply.

2. Always ask how bandwidth recommendations were arrived at?

There is a good chance that they were derived from traffic rate measurements with poor measurement granularity. At a minimum, make sure traffic rates are measured at a level comparable to the latency objective, e.g. if you are concerned with latency levels of 1 ms then traffic rates need to be measured with at least a granularity of 1ms.

No matter how knowledgeable your provider seems to be on these matters, our experience suggests that you should always verify the claims.

3. Don't forget about loss

Assuring low-loss or zero loss on market data networks is equally important to assuring low-latency. Remember to feed latency problems with bandwidth and feed loss problems with extra buffering.

4. Monitor Continuously

Experience tells us that conditions change frequently. For example average and peak message rates change as a result of market activity. This can cause higher than specified levels of latency and loss. It is important that these changes are monitored continuously and suitable alarms are triggered.

5. Verify

No matter how knowledgeable your provider seems to be on these matters, our experience suggests that you should always verify the claims. You may be surprised with the results.

Further Reading

To understand these issues in greater detail and to see how to measure, describe and specify bandwidth for low-latency market data in a manner that avoids the above pitfalls please read the following Corvil white paper entitled: "*Managing Performance in Financial Trading Networks*" by Fergal Toomey – Corvil Chief Scientist.

Acknowledgment

We would like to acknowledge the cooperation of BATS in the publication of this issue of Corvil Insight. It is much appreciated.



Corvil Ltd, 6 George's Dock, IFSC, Dublin 1, Ireland

T +353 1 859 1000 **Tech Support** +353 1 859 1010 E info@corvil.com W www.corvil.com

Copyright © 2009 Corvil Ltd. Corvil is a registered trademark of Corvil Ltd.
All other brand or product names are trademarks of their respective holders.