



PERIBIT PFA & SR-80 ANNOUNCEMENTS

FREQUENTLY ASKED QUESTIONS

Q1. Historically, Peribit has positioned itself as the “leader in instant bandwidth capacity”. With today’s PFA and SR-80 announcements, it appears the company is putting much more of an emphasis on “application performance.” Is there a shift in Peribit’s core competency/positioning as a result of today’s announcements?

A1. Not at all. Peribit continues to be the front runner in the bandwidth optimization market because of its unique ability to provide “instant WAN capacity”. Peribit originally discovered that it could apply DNA pattern matching algorithms to data networks to remove massive amounts of repetitive data from WAN links. The result was Peribit’s patent-pending Molecular Sequence Reduction (MSR) technology that instantly increases the capacity of existing networks by up to 10 times.

There are actually three common impediments to application performance: bandwidth, contention, and latency. Our core *MSR technology* addresses bandwidth limitations. Last year, Peribit added *essential QoS* to address application contention, enabling network managers to dynamically allocate this newfound bandwidth for specific WAN applications. Peribit’s new Packet Flow Acceleration (PFA) technology ensures that data networks are no longer susceptible to poor application performance caused by WAN latency.

By optimizing, allocating and prioritizing traffic, network managers can ensure maximum application performance. With today’s PFA announcement, we are addressing the third WAN application performance challenge: *latency*.

Q2. What is fueling customer demand for the new flagship platform, the SR-80?

A2. Peribit customers are deploying Sequence Reducers (SRs) in increasing numbers. More and more customers with initial deployments targeting 5-10 locations are immediately realizing the dramatic benefits Peribit SRs deliver in terms of increased network capacity, the elimination of application contention, and significantly improved WAN application performance. As a result, even though we are adding record number of new customers every quarter, nearly 40 percent of our orders each quarter come from Peribit's installed base.

Because of a rapid increase in the number of customers deploying over 120 SRs across their WANs (the maximum number of end point devices supported by a single SR-50/55 system), we're responding with the release of the SR-80. The SR-80 allows customers to deploy one device at their hub location to communicate with up to 320 other sequence reducers. This allows Peribit's customers to deploy hundreds of remote SRs with a smaller number of central site SR units, and it makes feasible the deployment of thousand-node and greater Peribit SR networks. In addition, the SR-80 platforms allow choice of either fiber or copper gigabit insertion at full wire speed with the increased reliability of hot swappable dual redundant power supplies.

Q3. Over 10% of Peribit's customer base are Fortune 1000 /Global 500 companies. Why will the SR-80 enable Peribit to make further in-roads into these large enterprise accounts?

A3. These large globally distributed organizations have WAN connections to hundreds and even thousands of locations. The SR-80 allows these customers to more easily scale their network without consuming valuable data center real estate at their core network locations or sacrificing central site reliability. Additionally, the vast majority of F1000 and global enterprises maintain network environments with very high data rates that require gigabit connectivity at the core of their networks. The SR-80 accommodates all of these requirements.

Q4. Peribit makes dramatic claims regarding the SR-80's ability to expand existing WAN link capacity up to ten times. What is the basis for this claim?

A4. The SR-80 utilizes an updated version of the same operating software as its predecessors, which includes our breakthrough MSR technology. MSR is the key enabling technology that eliminates repetitive WAN traffic in order to boost the bandwidth capacity of existing networks by up to ten times. These results are substantiated by actual customer deployments in over 200 networks and independent verification by The Tolly Group and Network Computing.

Q5. Peribit, like many vendors, is making a 'fast ROI' argument for its solutions. How does Peribit help customers develop real-world measurable ROI with its Sequence Reducer product family?

A5. Our ROI is very direct and measurable. In fact we call it a 'hard dollar' ROI because our customers can stop paying carriers incremental monthly charges for bandwidth, since we help them eliminate network link upgrades and adhere to minimum billable CIRs.

Typically, customer network traffic grows as a result of a new application rollouts, addition of more users, or a relocation of servers, data centers and offices. Prior to the availability of Peribit's solutions, customers facing these scenarios had to consider purchasing more capacity from their carrier, significantly increasing their monthly WAN link expenses, or reduce the level of performance for some applications to protect the others. Peribit is an alternative to this traditional scenario. Customers purchase Peribit Sequence Reducers as a one-time expense and effectively increase their link capacity by up to 10 times. When compared to purchasing this additional capacity, Peribit's Sequence Reducers typically provide a hard dollar ROI in well under 12 months – many times in less than 6 months.

On top of this hard dollar ROI, customers benefit from our QoS capabilities that enable them to leverage their 'new-found' network capacity in order to allocate and prioritize bandwidth for specific high-priority WAN applications. And with the release of today's PFA technology, Peribit customers will be able to realize considerable application performance improvements to significantly increase overall productivity across the enterprise for all applications.

Q6. The SR-80 is a substantial step up from the SR-50/55 platform. Since the SR-80 is clearly a central site solution, where will the SR-50/55 be deployed?

A6. The SR-50 and SR-55 will continue to play an important role in our product line for those customers deploying fewer than 120 endpoint devices in both central site and remote office locations. The SR-50 and SR-55 platforms, which remain viable central site platforms, are priced and configured to meet the needs of customers that do not forecast growth beyond these parameters. In addition, with the introduction of SR-80s at centralized locations, SR-50 or SR-55 platforms currently located in central sites can easily be redeployed into remote locations as the customer's rollout of Peribit solutions grows.

Q7. The new PFA technology claims to slash WAN latency. What exactly are the causes of WAN latency?

A7. WAN latency is caused by several factors or sources. In private networks the primary cause of WAN latency is serialization and propagation delay, usually described in milliseconds of round-trip-time or "RTT(ms)". This is the amount of time it takes for data to be placed on the network at its given link connection-speed, and the amount of time the signals take to reach the destination at the speed of light.

Best case RTT(ms) scenarios for metropolitan private networks range from 10-50ms, for national networks from 100-200ms, and for intercontinental links from 200-500ms. In addition, for IP VPNs, or more complex internally routed enterprise networks, the number of router hops across a network can significantly increase the end to end latency. On average, a router hop introduces from 50-200 additional milliseconds of latency. Therefore VPN connections or private WANs that have several hops can easily introduce over 500 ms of latency.

Finally, due to the long-distances of extra-terrestrial radio links, latency becomes extremely high on satellite links, typically adding 275 ms or more per satellite hop to the above numbers. On multi-router hop satellite links between the most distant locations in the world true end-to-end latency has been measured at up to 2270 RTT(ms) on the public internet (source: University of Wisconsin).

In many cases, customers are completely unaware of the impact WAN latency has on enterprise-wide application performance. However, those customers deploying Peribit Sequence Reducers can easily optimize all aspects of WAN performance to maximize the performance of their applications. With Peribit's products customers can easily determine where WAN latency is occurring in order to proactively resolve link delays across their networks.

Peribit's Sequence Reducers with PFA dynamically and automatically repair the damage done to application performance by latency.

Q8. How does WAN latency impact an enterprise's bottom line? Please provide real-world examples.

A8. Latency delays acknowledgements between send and receive endpoints and this impedes an application's ability to send additional data while it's in a "wait state". This can result in sluggish screen response times with transactional 'short-flow' applications such as ERP or CRM packaged apps (SAP, Oracle, JDEdwards), Web-services based apps (Peoplesoft, Microsoft Apps), etc.

Latency can also dramatically bog down database lookups, which can have serious consequences for customer service applications. And, it can significantly impact overnight data backup, downloading, and replication throughput, which can narrow the windows of time during the business day that enterprises can service customers, as most applications have to perform this mission-critical function between business cycles.

Q9. Data replication and distributed caching have been deployed for some time to address WAN latency. Why is PFA a superior approach?

A9. It is a well-established fact that caching does not work for many enterprise applications. Caching is typically focused on the diminishing components of static web content on enterprise networks, which is why it has not been widely successful in enterprise networks. Data replication solutions like Network Appliance's Snap Mirror actually drive more data over the WAN, heightening the need for Peribit solutions to accelerate this replication processing over the WAN. In fact, Peribit customers have deployed Sequence Reducers to support data replication solutions like this around the world.

Q10. Peribit is claiming its PFA technology can reduce application response times and increase WAN throughput by 'up to 4 times'. What metrics is Peribit using to substantiate these claims?

A10. These claims apply to applications which are capacity constrained as a result of latency, which is an often ignored and unidentified problem. To clarify how PFA improves application response time and throughput lets look at a specific example:

In an enterprise network a 100 megabyte file is transmitted over a WAN link at 1 Mbps with 40 milliseconds of latency. Transmitting this file normally

takes 14 minutes. When we insert MSR the transmission time is reduced to about 9 minutes. It should have completed in less than 4 minutes, because MSR provided a 75% reduction rate (or 4X capacity gain). The benefit of MSR comes from reducing the amount of data sent across this link. This application left over half of MSR's new-found capacity unused, albeit available for other applications.

In this example latency is slowing down the sender while it waits for acknowledgements from the receiver, and it simply can not run any faster. However, when we add PFA to MSR the transmission completes in 3 minutes, 56 seconds. This demonstrates a 2.5 times improvement in the transmission time over MSR alone. PFA can produce 2-4x improvement depending on the latency and other link conditions and allows networks to deliver maximum WAN application performance.

Q11. Is the effect of PFA additive? Does it further reduce response times by 2-4X over the reduction enabled by MSR?

A11. Not exactly. PFA will have a dramatic affect on overall application performance for applications which are latency constrained. MSR eases the network capacity constraints on WANs for all applications, which is the first order problem faced by enterprises today and produces dramatic and rapid ROIs. Essential QoS addresses congestion by effectively allocating this new-found bandwidth and prioritizing key applications. Once capacity and congestion constraints are removed, however, the impact of latency must be mitigated in order to maximize application performance. Therefore, while MSR can increase the capacity of a link by up to 10 times, PFA can ensure that this capacity is available to the applications that need it.

Q12. Explain the relationship between MSR and PFA?

A.12. While MSR is used to resolve network bandwidth constraints, it also dynamically detects whenever latency becomes an issue for any individual traffic flow. If it is determined latency is an impediment to optimal performance, PFA will begin to operate on the specific flow to mitigate the inefficiency caused, thereby reducing application response time and enabling more rapid data transfers.

Q13. How exactly does PFA speed application response time?

- A13. PFA dramatically speeds application response time by eliminating gaps in TCP packet flows that cause poor throughput and reducing inefficiency in the connection setup and processing between end systems over the WAN that cause slow response time. These gaps and inefficiencies are caused by delays in protocol-level acknowledgments of previous transmissions.

Peribit's Packet Flow Acceleration (PFA) reduces the impact of latency on enterprise applications by significantly reducing the amount of idle time transmitting systems wait for delayed packet acknowledgements, thus keeping more data in flight and eliminating the gaps that cause inefficiency. PFA increases the responsiveness and throughput of latency constrained applications up to four times, enabling higher employee productivity and allowing bandwidth-intensive applications such as data replication or backups to complete within target windows.

PFA is a dramatic departure from TCP "spoofing" technology that attempts to improve TCP performance by altering latency effects with new protocols. Spoofing terminates TCP sessions on the transmit and receive ends by pretending to be (proxying) the remote TCP end-points. It "spoofs" TCP hosts by acknowledging data before it has even been sent across the WAN, and takes over responsibility for the end to end reliable communication and congestion management using other protocols. In addition, TCP spoofing suffers from poor scalability since every communication session end point is replicated inside the equipment for both inbound and outbound sessions. As a result, spoofing solutions are not widely deployed today.

In contrast, PFA is fully transparent to TCP host communications and does not maintain session state as a 'proxy' to the hosts. What PFA does is accelerate TCP by intelligently allowing larger windows of data to be transmitted over the WAN. As a result, TCP takes care of all reliable end to end transport functions and it's scalability remains unaffected. PFA never holds or delays acknowledgments, does not terminate TCP or interfere with congestion control, and does not require any additional congestion control or reliability mechanisms. In fact, Peribit's approach has the support of industry luminaries as the ideal way to counteract the impact of latency.

"PFA is an innovative approach that leverages the company's unique position at both ends of a WAN link," said Dr. Vern Paxson, senior scientist with the ICSI Center for Internet Research. "People want data networks to deliver the same immediacy whether they're connecting across the street or across the world. Peribit transparently optimizes TCP communications, while still taking care to honor TCP's congestion control mechanisms." The International Computer Science Institute (ICSI) is an independent non-profit

research institute, located in Berkeley, California and closely affiliated with the University of California, Berkeley.

In summary, PFA is highly transparent, reliable, and scalable since it maintains a minimal internal TCP state and only operates on flows that become limited by latency. It also avoids adding complexity in the network because TCP operates end-to-end as the standard congestion-control mechanism. Equally important, no parameter tuning or ongoing management is required to take advantage of PFA since it is tightly integrated with MSR in the Sequencer Reducer operating software.

Q14. Is PFA a standalone technology, or does it interoperate with Peribit's Molecular Sequence Reduction (MSR) technology? How do customers "purchase" this new functionality?

A14. PFA is bundled into Peribit's Sequence Reducers at no additional charge. It is integrated with MSR in SRS OS 4.0 and later, and delivered as a packaged solution for the customer. There are no additional configurations required to add this capability to a customer's network.

Q15. PFA sounds like TCP *rate shaping*. How does PFA differ from this approach, and how does it avoid "breaking" normal TCP/IP behavior?

A15. PFA is dramatically different than TCP rate shaping. TCP rate shaping, also sometimes called *ack pacing*, delays and substitutes delivery of acknowledgements in order to slow down targeted TCP connections, a method that can cause serious reliability issues with TCP connectivity if connections fail or packets are lost. TCP rate shaping individually manages each TCP acknowledgement to slow down or 'pace' lower priority flows and to allow other high-priority flows to run faster. PFA is designed to overcome the impact of latency on all TCP flows without interfering with or delaying delivery of acknowledgements. PFA allows more data to cross the network by eliminating the inefficiencies in TCP flows caused by latency.

Q16. Is PFA a standards-based approach to improving application response times?

A16. PFA is Peribit's technology for optimizing standards-based TCP/IP packet flows in a fully standards-compatible way. PFA does not change the end to end communication patterns between senders and receivers and generates no protocols itself. The semantics of TCP flows are not changed when PFA-enabled Sequence Reducers are sitting between either end of a TCP session. PFA simply and effectively improves the performance of applications by

monitoring and dynamically tuning the advertised receiver window sizes in normal TCP communication whenever it detects that latency is causing inefficient use of bandwidth by connections into that receiver. Because of its location close to each end of the WAN connection, where most of the latency is introduced, PFA can efficiently detect which flows are not optimized and correct this inefficiency.

Q17. Does PFA create any installation or configuration complexities for network administrators?

A17. PFA is completely transparent from the perspective of network administrators. PFA can be enabled for all applications or only specific applications on each SR. The SR devices automatically determine which of the selected TCP flows could be enhanced with PFA by looking for windows that have already reached the maximum number of packets in flight and whose acknowledgments have not yet been received by the SR on the senders side. PFA is then automatically applied to these flows to result in better response time and throughput as a result of increased efficiency.

###