

White Paper

Considerations for Using T.38 versus G.711 for Fax over IP

Considerations for Using T.38 versus G.711 for Fax over IP

White Paper

Executive Summary

Businesses migrating to Voice over IP (VoIP) often find it desirable to move their fax traffic onto the IP network as well. However, VoIP networks are, as the name implies, optimized for voice traffic, and businesses implementing a Fax over IP (FoIP) solution as part of their fax system can benefit from understanding their options for FoIP transport methods.

This white paper compares the performance of the two principal options for sending faxes over an IP network: T.38 fax relay and G.711 fax pass-through. The V.17 and V.34 modem standards are briefly discussed and also used for comparison when used with T.38 and G.711 for performance against IP network impairments, such as latency, packet loss, and jitter. The impact of these network impairments on call control, fax control, and image data are described, as well as the impact of the frequency of these impairments. General considerations for network performance are also covered, as are sections describing how metrics within the Dialogic® Brooktrout® Bfv API can be used to detect IP network impairments.

Table of Contents

Introduction	2
FoIP Transport Methods	2
T.38 versus G.711	2
V.17 versus V.34	3
Redundancy	3
ECM	3
Typical IP Network Impairments	4
Latency	4
Packet Loss	4
Jitter	4
Impact of Network Impairments	5
Call Control	5
Fax Control	5
Image Data	6
Impairment Frequency	6
Observed Impacts of Network Impairment	6
Effects of Single Packet loss	6
Considerations for Network Performance	12
Dialogic® Brooktrout® SR140 Fax Software	14
Appendix A—Detecting Network Issues using Fax Quality Metrics	15
Appendix B—Test Methodology	19
References	19
For More Information	19

Introduction

In the 1990s, “fax” and “T.30” were so entwined that they were basically synonymous. T.30, as the ITU recommendation for transmitting faxes over the general switched telephone network, includes mechanisms to handle background noise and spikes of interference on the telephone line. For example, poor signal quality can be accommodated by lowering the transmission speed, and spikes of noise can be handled by retrying any operations that are lost during the spike.

But T.30 was not designed to deal with IP network impairments such as packet loss, which can result in large gaps in fax data and cannot be eliminated by a lower transmission speed. In fact, lower speeds may actually make the performance worse by generating additional network traffic that will be exposed to packet loss. And retries may not always recover missed data when packet loss is frequent, because the retries can suffer packet loss as well. Even the retry requests themselves may be lost.

When upgrading from traditional analog fax to a Fax over IP (FoIP) solution, it is necessary to provide a sufficient level of network performance to support reliable operation. Fax transport over an IP network can be provided by T.38 fax relay and G.711 pass-through, each one having different network requirements for reliable operation.

This white paper describes and compares the T.38 and G.711 fax transport methods for FoIP, using both V.17 and V.34 modem standards, for performance against typical IP network impairments, such as latency, packet loss, and jitter. It then discusses the impact of these network impairments on call control, fax control, and image data, as well as the impact of the frequency of these impairments. This paper also provides considerations for network performance for each fax transport method, and includes information pertaining to the ways metrics within the Dialogic® Brooktrout® Bfv API can be used to detect IP network impairments.

FoIP Transport Methods

When implementing a FoIP solution, the transport method used (T.38 or G.711), in addition to several fax settings such as fax speed, redundancy, and error correction, can have a significant impact on FoIP performance for common network impairments.

T.38 versus G.711

T.38 fax relay is an ITU-T recommendation that allows for fax data to be carried over IP networks. Data is transmitted directly in T.38 without being converted to an audio stream, and results in a significant reduction in the bandwidth needed. T.38 also supports data and controls redundancy to mitigate the effects of packet loss.

Some disadvantages of T.38 are that gateway support for fax parameters, such as V.34 transmission speeds and Error Correction Mode (ECM), is not universal. Also, in the current mixed network environment of packet-based and circuit-switched (that is, PSTN) connections, T.38 often has a transcoding overhead. That can add latency and cost to fax services.

G.711 is an ITU-T recommendation for Pulse Code Modulation (PCM) of voice frequencies. It uses an uncompressed format and requires high bandwidth, typically about 64 kbps. Using G.711 as the transport method for FoIP is an extension of traditional PSTN audio-based faxing. The digital fax data is converted to a PCM audio stream and then sent as G.711 Real-time Transport Protocol (RTP) packets.

G.711 has not been optimized for fax transport over IP networks, and does not typically support packet redundancy. Having been developed for voice, G.711 allows the transmission of missing audio because any gaps would be filled in by a human listener. But when used to transmit modem data, any loss of packets is significant, because the receiver has no way to recreate the missing data.

Some packet networks may not be fax-aware, and may optimize the G.711 stream for voice with the use of silence suppression, echo cancellation, or transcoding to a higher compression codec. Such optimizations can cause a loss of data and prevent FoIP from operating. This may force the use of a dedicated G.711 fax trunk to provide reliable fax performance. However, G.711 is an inherently simpler approach to fax than T.38, so interoperability issues between different vendors' products is less common with G.711 than may be encountered with T.38. The cost for a G.711 approach may also be lower than T.38 because it can leverage voice data infrastructure.

V.17 versus V.34

V.17 and V.34 are two commonly used modem specifications for fax that were developed in the PSTN era of faxing, which explains why neither was designed with IP network impairments in mind. V.17 modulation supports bitrates up to 14400. The final transmission speed that is used is determined during a short training cycle, when a known pattern of bits is transmitted to see if it is received successfully. The modem will try progressively lower speeds until the training pattern is correctly received (known as "training down"). The V.34 specification was introduced later to achieve higher transmission speeds over traditional PSTN phone lines, and supports bitrates up to 33600. V.34 uses a longer and more complex training cycle than V.17 to determine the maximum speed the line can support. The comparatively higher complexity of V.34 can make it more vulnerable than V.17 to speed train downs when there are network impairments such as packet loss.

Redundancy

Network impairments typically result in lost IP packets. Redundancy is a method whereby transmitted information is replicated and repeated in several packets. By repeating data in this way, the probability is higher that in the end all of the transmitted information will reach the receiving side, even if a few packets are lost. This can also reduce the need to re-transmit missed information, which in turn may reduce the transmit time for a fax.

T.38 generally supports two types of redundancy:

- **Control**—Refers to IP packets that contain fax control commands. This is sometimes referred to as "low speed" redundancy because in traditional analog fax, these commands are transmitted at a low data rate of 300 or 1200 bps.
- **Data**—Refers to IP packets that contain fax image data, and is sometimes referred to as "high-speed" redundancy.

G.711 can also support redundancy through the use of redundant RTP, but support for this is not currently widespread.

ECM

Error Correction Mode (ECM) is a traditional fax checksumming method applied to blocks of fax data. Not all fax devices support it, but for those that do, each segment of fax data is sent with a checksum that is verified by the receiving side. If part of the data is corrupted or missing, the receiving side will request that it be re-transmitted.

When ECM is not used, missing data will simply be omitted from the received fax image, causing some degradation in image quality. Because network impairments typically result in the loss of data, ECM can help to preserve image quality in these situations. But ECM usage can cause an increase in both the transmission time and the number of unsuccessful faxes (that is, fax attempts that are not completed). The transmission time can increase because missing data needs to be requested and re-transmitted. The number of unsuccessful faxes can increase because impairments may corrupt re-transmissions, causing the fax to be aborted when the maximum number of retries is reached.

V.34 requires the use of ECM, but with V.17 the use of ECM is optional and often configurable to be on or off.

Typical IP Network Impairments

The two most common IP network impairments that impact FoIP are latency and packet loss.

Latency

Latency refers to the amount of time it takes transmitted data to reach its destination. Round trip latency refers to the amount of time it takes transmitted data to reach its destination plus the time for the destination's response to be returned.

As IP data is sent from a fax endpoint to a receiving endpoint, the data packets are relayed through a series of network elements. The first would normally be a series of Ethernet routers. Then a gateway may receive the data and transcode it into another format. For example, a gateway may transcode a T.38 digital fax into a PSTN audio stream. In some cases, the fax may even be routed through more than one gateway.

Every network element adds some amount of latency to the data. Some add a very small amount on the order of milliseconds, whereas others can add significantly more time. If the data travels over the traditional PSTN network, it may also add delays depending on what type of transmission devices are used to reach the destination.

In general, the amount of latency for a given fax call can vary widely, depending on such factors as the source and destination locations, as well as what network elements are between them.

Packet Loss

Packet loss can occur when network congestion is high and a network element is unable to relay all the packets it receives. These congestion periods may be short lived, such as when a large data file is being transferred on a network. This can also occur when high priority data, such as real time voice data, is given priority over other network traffic. Hardware causes, such as poor signal quality on a cable, can also lead to packet loss.

Burst Packet Loss

Burst packet loss refers to a series of consecutive IP packets from a stream not reaching its destination. The length of the burst often determines how negatively it will affect a FoIP call.

Single Packet Loss

Single packet loss refers to the occasional loss of single packets (that is, non-consecutive packets) from a data stream. This type of loss is not as significant as burst losses, but can cause serious issues if it occurs frequently.

Jitter

Jitter involves elements of both latency and packet loss. Jitter refers to variations in latency the network adds to transmitted packets. Because faxes are sent in real time, they require a continuous stream of data to transmit successfully, particularly if the data is being carried as an audio stream. If a packet of data has not been received when it is needed, then the net effect is the same as a lost packet. If several packets are received too late to be used, then that has the same net effect as a burst of lost packets. Even if the late packets eventually arrive, it will be too late for the receiving side to use them, and thus they will be discarded.

For FoIP, packet loss due to jitter can be well controlled with the use of jitter buffers. This process buffers several packets before they are read so that the receiver will not run out of packets if/when some are delayed. But a drawback of jitter buffers is that they add latency. And when multiple network elements are in a connection employing jitter buffers, then the latency effect will be additive.

Another source of packet loss related to jitter can occur when G.711 faxes experience clock skews. This happens when the sender and receiver are using unsynchronized clock speeds for the audio stream. If the sender is creating packets slightly faster than the receiver is reading them, eventually the jitter buffer of the receiver will overflow and a packet will be lost. Or if the receiver is reading packets faster than they are being sent, eventually the jitter buffer will run dry, and the receiver will not have a packet to read when it needs it.

So although jitter is a form of latency, its impact on fax performance is usually related to how much packet loss it causes, rather than the latency itself.

Impact of Network Impairments

Network impairments can cause faxes to fail in many different ways. Several layers and phases are employed by FoIP, and each one can react differently when impacted by network impairment. The specific reaction is largely determined by the severity of the impairment and which phase of the call is in progress when the impairment happens. On an application level, fatal errors will typically be reported by the fax API function that was in progress when the impairment occurred. Non-fatal errors are not reported by all types of network elements, but can be observed by products that support the collection of network quality metrics, such as those in the Dialogic® Brooktrout® Bfv API, the application programming interface used for developing fax solutions based on both the Brooktrout Dialogic® Brooktrout® SR140 Fax Software and the Dialogic® Brooktrout® Fax Boards (see “Detecting Network Issues using Fax Quality Metrics” in Appendix A).

Call Control

The first phase of a fax that can be impacted by network impairments is the call control phase. For FoIP, call connections are generally controlled with the Session Initiation Protocol (SIP) or H.323 protocols. For example, if packets are lost on the network and the call setup packet is dropped, then the receiving side will not know that a call request has occurred. The sender will wait for a response, eventually time out, and end with a “no answer” or “line busy” error. If the response packet to the call setup is dropped, the result would be similar. To complete the call, the fax server or user would need to make another attempt to send the fax.

Other call control operations that can be impacted include the fax transport negotiations (that is, setting up T.38 or G.711) and call termination.

Fax Control

Once a call has successfully transitioned into fax mode, using either T.38 or G.711, the call enters the fax control phase. During this time, traditional T.30 fax operations take place, such as training the modem to the highest working bitrate, and negotiating to send the next page in the fax. Packet loss during this time could cause the fax speed to train to a lower speed, or cause the fax to fail altogether. Lower speeds will increase the fax transmission time.

If fax control commands or their acknowledgements are dropped, it triggers several retries (usually up to three) for those commands until they succeed, or until one side gives up and ends the call. The result could be that all or just part of the fax is truncated at the receiving end.

High latency may cause fax control failures because T.30 has required timer values for completing some command sequences. For example, most commands use a three second timer for receiving a response. If a response is excessively delayed, these timers will be exceeded and the call will fail. Fax calls that end prematurely will need to be re-sent by the fax server or a user.

Image Data

During the fax data phase, the actual image data for each fax page is transmitted. Depending on whether ECM is being used, missed data will be re-requested or simply left out of the received image. If re-requested, the fax transmission time will increase. If packets for the re-transmission requests are lost, then there will be further delays and possibly failure of the fax if retry limits are reached. If ECM is not used, then lost data will result in omissions in the received fax image, usually in the form of missing or repeated horizontal strips.

Impairment Frequency

How often the network impairment occurs can also have a large impact on the overall fax completion rate. If the impairment is continuous, (for example, consistently dropping some percentage of transmitted packets) then the impact on the fax completion rate will be more pronounced. Also, with a continuous impairment, faxes with longer durations (that is, those faxes with a large number of pages), will have a greater exposure to the impairment, thus increasing the chance the fax will fail or terminate prematurely.

If the impairment is infrequent and short in duration, it will have less of an overall impact on fax performance. However, it is important to note that some packets within a fax transmission are more critical and vulnerable to loss than others. Data packets, for example, can often be lost without causing the fax call to fail, but the loss of certain control packets can more readily lead to a failed call. With that in mind, even a small and infrequent network issue could result in occasional fax failures, if, by chance, it impacts a particularly important set of packets.

Observed Impacts of Network Impairment

Figures 1 through 12 are test examples showing the effect of network impairments on different FoIP transport configurations. For these tests (see Appendix B for more information regarding test methodology), two servers, each equipped with a FoIP test application created using Brooktrout SR140, were connected using an artificially impaired network:

Brooktrout SR140 ← → Artificially impaired IP network ← → Brooktrout SR140

Effects of Single Packet loss

As a general consideration, Table 1 shows the levels of single packet loss that can be observed on various types of networks [ITU-T].

Network Type	Single Packet Loss
Well managed IP network (that is, supports real-time applications with strict constraints)	< 0.05%
Partially managed IP network	< 2%
Unmanaged IP network (Internet)	< 20%

Table 1. Single Packet Loss on Various Network Types

Figure 1 shows an example of how T.38 used with packet redundancy can provide excellent protection against single packet loss. As is also shown, when redundancy was disabled, the fax transmissions were more vulnerable to failure from the packet loss.




Figure 1. T.38 Fax Completion Rate versus Packet Loss

Figure 2 shows an example in which G.711 was very vulnerable to single packet loss, whereby even a small percentage of single packet loss causes unacceptable failure rates. Disabling ECM was shown to increase the fax completion rate, but came at the expense of reduced image quality of the received fax.




Figure 2. G.711 Fax Completion Rate versus Packet Loss

When comparing Figures 1 and 2, note that T.38 with redundancy was vastly superior to G.711 at dealing with single packet loss. Even with redundancy disabled, T.38 was still considerably more robust against packet loss than G.711.

Effect of Fax Length with Single Packet Loss

Figures 3 and 4 are examples of when the network impairment is held constant while the number of fax pages sent is increased from one to nine; these Figures also show the fax length impact on call duration and the fax success rate.

In Figure 3, note the effect of a 0.25% single packet loss on fax transmission time. For G.711, this amount of impairment had a significant impact and caused the duration of the call to increase significantly. Also for G.711, as the number of pages increased, more packets were exposed to the network impairment, which magnified the effect of the impairment. The elongation was primarily due to command retries and the retransmission of data.




Figure 3. Call Duration versus Number of Pages

By contrast, the faxes sent using T.38 were virtually unaffected by the small amount of packet loss and the call duration increased smoothly to transmit the additional pages with no retries needed. The T.38/V.17 configuration represented the pages being transmitted at the maximum 14400 bps speed, whereas the T.38/V.34 configuration represented the pages being transmitted at the maximum 33600 bps speed, giving the same results as if there were no impairment.

As with call duration, note in Figure 4 that the longer the G.711 fax data stream was exposed to the impairment, the more damage it did to the success rate. A minor impairment that might rarely cause issues for a one-page G.711 fax may significantly impact the ability to send a multi-page fax by causing it to end prematurely. In the Figure 4 test example, a T.38 fax was able to maintain a 100% success rate regardless of its number of pages.




Figure 4. Fax Completion Rate versus Number of Pages

Effects of Burst Packet Loss

As a general consideration, Table 2 shows the levels of burst loss that might be observed on various types of networks [ITU-T].

Network Type	Burst Packet Loss
Well managed IP network (that is, supports real-time applications with strict constraints)	Random loss only
Partially managed IP network	40-200 ms
Unmanaged IP network (Internet)	40-10,000 ms

Table 2. Burst Packet Loss on Various Network Types

For the burst loss tests (see the example in Figure 5), the network traffic was dropped for increasing lengths of time during each one-second interval. At shorter burst lengths, this resulted primarily in single packets lost, while with longer bursts there were several consecutive packets lost. For G.711 FoIP, packets are typically sent at fixed intervals of 20 ms; for T.38, the packet transmission rate is not fixed and will vary as the call progresses through different phases of call setup, parameter negotiation, and image transmission.

During the testing, burst loss with T.38 caused the success rate to drop more sharply than with single packet loss because losing consecutive packets was more damaging to control operations than losing single packets. It also reduced the effectiveness of packet redundancy. Redundancy works by replicating data across sequential packets, so when several sequential packets are

lost, it increases the odds of losing all of the redundant copies for a particular packet. In general, V.34 tends to be more sensitive to packet loss than V.17 due to its longer training cycle and higher transmission speed.




Figure 5. T.38 Fax Completion Rate versus Burst Packet Loss

As with single packet loss, Figure 6 is an example showing that G.711 fax is very sensitive to any packet loss. Burst loss of very short duration was sufficient to greatly reduce the fax success rate.




Figure 6. G.711 Fax Completion Rate versus Burst Packet Loss

In comparing T.38 and G.711 performance under burst loss (Figures 5 and 6), note once again that T.38 was far more robust, even without redundancy enabled. At a level of 30 ms burst loss per second, all of the G.711 faxes were failing, while the T.38 faxes continued to have a relatively high completion rate.

Effect of Latency

As a general consideration, Table 3 shows the amount of round trip latency that may be observed on various types of networks [ITU-T]. Note that this table does not factor in latency that may be added by transcoding, jitter buffers, and/or processing time at the remote endpoint for generating responses.

Network Type	Network Round Trip Latency (not including transcoding, jitter buffers, or response generation)
Well managed IP network (that is, supports real time applications with strict constraints)	40 to 200 ms (regional)
	180 to 600 ms (intercontinental)
Partially managed IP network	100 to 200 ms (regional)
	180 to 800 ms (intercontinental)
Unmanaged IP network (Internet)	100 to 1000 ms

Table 3. Round Trip Latency on Various Network Types

With increasing latency on the network connection, T.38 and G.711 performed similarly, with G.711 being slightly less tolerant of latency than T.38 (see Figure 7). This was due to the extra latency added by the G.711 jitter buffer on both the sending and receiving sides of a call. The use of redundancy and ECM had little impact on tolerance to latency.




Figure 7. Fax Completion Rate versus Round trip Latency

FoIP was generally tolerant of network latency up to about 2.5 seconds, beyond which all fax attempts began to fail. This was due to timer restrictions imposed by T.30. Most commands have a three second retry timeout, so when responses are delayed beyond that, the fax endpoints will decide the other endpoint is not responding, and terminate the call.

In an actual deployment, G.711 might experience less network latency than T.38, because gateways can add latency when transcoding from T.38 to PSTN connections. There might also be cases where T.38 calls pass through multiple gateways and are transcoded more than once. This can be mitigated somewhat by fax spoofing that is done by most gateways. With fax spoofing, the gateway will send “keep alive” messages on the PSTN side of a transcoded FoIP call, to prevent a traditional PSTN fax device from timing out when there is high latency on the IP side of the call.

In T.30, the three second timeout is generally referred to as the “Timer 4” or T4 timeout. For some products, such as the Brooktrout SR140 that was used for this testing, the setting for the T4 timeout can be manually configured. In Figure 8, the T4 timeout was increased to a value of four seconds.

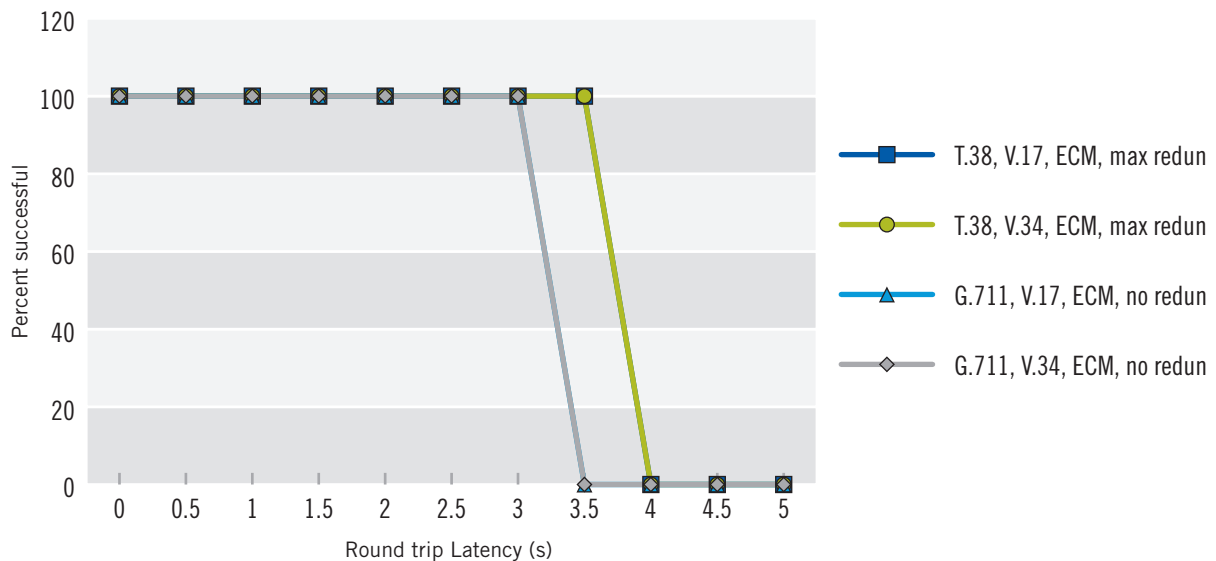


Figure 8. Fax Completion Rate with T4 Set to 4 Seconds

As was expected, changing the T4 timer to 4 seconds increased the tolerance for latency by one second. Thus, such a change might be useful in networks that have very high latency. But note that because modifying the T4 timer above three seconds places it outside of normal T.30 specifications, doing so may not be compatible with all fax devices.

Considerations for Network Performance

In terms of robustness, T.38 fax relay was shown in these examples to have a much higher tolerance for network impairments than G.711 pass-through. However, a network properly optimized for G.711 FoIP can also provide a reliable fax solution.

The goal for FoIP installations is to have zero failures attributable to the locally managed network environment. If measurable failures are due to network impairments, then sending faxes of long duration will magnify the errors, making it difficult to transmit such long duration faxes.

Table 4 provides a summary of the network impairments that might be expected for different types of networks.

Network type	Network Round Trip Latency (not including transcoding, jitter buffers, or response generation)	Single Packet Loss	Burst Packet Loss
Well managed IP network (that is, supports real time applications with strict constraints)	40 to 200 ms (regional)	< 0.05%	Random loss only
	180 to 600 ms (intercontinental)		
Partially managed IP network	100 to 200 ms (regional)	< 2%	40-200 ms
	180 to 800 ms (intercontinental)		
Unmanaged IP network (Internet)	100 to 1000 ms	< 20%	40-10,000 ms

Table 4. Network Impairments for Various Network Types

Table 5 shows levels of network errors that typically can be tolerated for commonly used FoIP transport methods. These numbers are approximate, but generally can provide acceptable fax performance in most FoIP environments. Results will vary depending on factors such as the complexity and length of faxes sent, the type of gateways used, and the type of endpoints involved. The transport methods in Table 5 are listed in the order of how robust they are against network impairments, from most robust to least robust.

Transport Method	Round Trip Latency Permitted	Single Packet Loss Permitted	Burst Loss Permitted	Recommended Network Type
T.38, V.17, ECM, with redundancy	< 2 seconds	< 20%	< 50 ms	Well managed or partially managed
T.38, V.34, ECM, with redundancy	< 2 seconds	< 5%	< 25 ms	Well managed or partially managed
T.38, V.17, ECM, no redundancy	< 2 seconds	< 1%	< 25 ms	Well managed or partially managed
G.711, V.17, ECM, no redundancy	< 2 seconds	None	None	Well managed only
G.711, V.34, ECM, no redundancy	< 2 seconds	None	None	Well managed only

Table 5. Typical Tolerance for Various Network Impairments

It is important to note that:

- To preserve image quality when using V.17, it is considered “best practice” to always enable ECM. Even though enabling ECM can slightly lower the success rate under high impairments and increase transmission time, loss of image data is rarely acceptable.
- When using T.38, redundancy should be configured to its maximum values. But note that this increases the amount of bandwidth used and might not be necessary in all cases.
- Jitter is not generally an issue for T.38, because the packets are not sent at regular intervals as they are for a G.711 audio stream. When using G.711, jitter buffers should be set to a fixed size that is greater than the amount of jitter present on the network. This will prevent the jitter from causing packet loss in the G.711 stream.
- When using G.711 for FoIP, consider optimizing the network to prioritize fax traffic to prevent packet loss. Network elements should be fax aware, such that voice optimizations such as silence compression are disabled and whereby jitter buffers are set sufficiently deep to preclude packet loss.

Dialogic® Brooktrout® SR140 Fax Software

Brooktrout SR140 fax software was used to collect the performance data for this paper. The Brooktrout SR140 software is a host-based FoIP engine that leverages field-proven Brooktrout® fax technology to deliver high levels of performance, reliability, and scalability that developers can use to create software for fax servers and other fax applications.

Brooktrout SR140 can deliver highly reliable FoIP by using T.38 to send real-time fax messages over IP networks. Because T.38 retains the standard T.30 fax data stream, it can be used with legacy T.30-based devices and with newer T.38 solutions. With over 20 years of experience, the Dialogic® Brooktrout® T.30 implementation can lay claim to being the most thoroughly field-tested offering. Dialogic has refined its T.30 implementation for numerous T.30 variants, which helps make fax processing reliable, saving users time and money.

G.711 fax pass-through is also supported by the Brooktrout SR140 software. While more susceptible to possible IP network issues, such as those described in this paper, G.711 fax pass-through provides an option for enabling FoIP when T.38 is not supported, or when the IP network meets the requirements for successful fax transport using G.711.

FoIP solutions built using Brooktrout SR140 fax software can send and receive IP faxes at up to 33.6 kbps, based on the V.34 fax standard. Not only can Brooktrout SR140 process fax at twice the speed of 14.4 kbps fax boards, but it also supports V.8 fast handshaking and advanced compression, which can cut call setup and session-management time by one-third. A document that can be faxed in one minute with a 14.4 kbps intelligent fax board can be faxed in less than 30 seconds with an application using Brooktrout SR140. This can translate into significant savings on long distance toll charges.

For more information on the Brooktrout SR140 fax software, visit http://www.dialogic.com/products/ip_enabled/FoIP/SR_140.htm.

Appendix A—Detecting Network Issues using Fax Quality Metrics

Based on ITU E.450 standards, the Brooktrout Bfv API has several fax quality metrics that allow an application to monitor network performance without needing external test equipment. The results shown in Figures 9 through 12 are based on testing conducted using the test methodology described in Appendix B and reported by the metrics within the Brooktrout Bfv API.

A good general indicator for network issues is the fax Figure of Merit (FOM) metric. The indicators for the FOM metric are measured on a scale of one to seven, with one indicating an error-free fax, and two up to seven indicating increasingly greater issues. FOM values consistently higher than one should be investigated, to determine and, if possible, eliminate the cause. Figure 9 shows the effect that increasing burst losses had on the FOM metric.

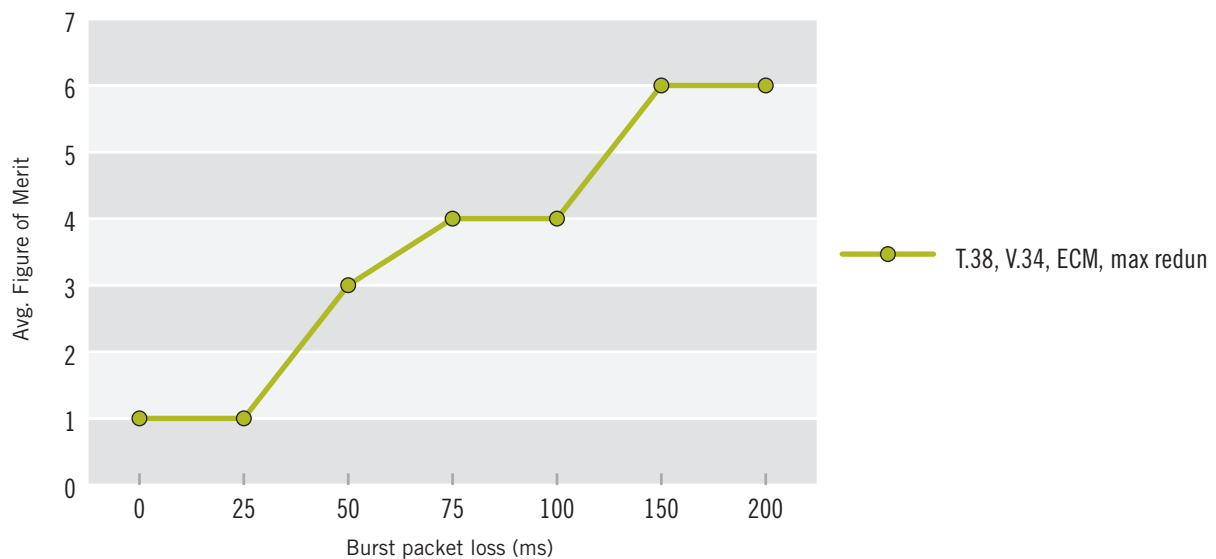


Figure 9. FOM versus Burst Packet Loss

When the FOM suggests there may be network issues, the cause of the issue generally can be narrowed down by examining more specific targeted metrics. For example, packet loss can be monitored for T.38 using the metrics shown in Table 6.

T.38 Metric	Value Being Measured
T38RcvrPackets	Number of packets received
T38RcvrLostPackets	Number of lost packets
T38RcvrLostPackets_NS	Number of single lost packets
T38RcvrLostPackets_N23	Number of 2 and 3 consecutive lost packets
T38RcvrLostPackets_N4	Number of 4 consecutive packets lost

Table 6. T.38 Metrics in the Brooktrout Bfv API

Burst loss will be seen in the N23 and N4 metrics, whereas single packet loss will be seen in the NS metric. An application could calculate the percentage of packet loss that is being seen and issue a warning when it exceeds the recommended guidelines for the transport mode being used.

Similarly, packet loss can be monitored for G.711 using the RTP metrics that are shown in Table 7.

RTP Metric	Value Being Measured
RTPRcvrLostPackets	Number of lost packets
RTPRcvrLostPackets_NS	Number of single lost packets
RTPRcvrLostPackets_N23	Number of 2 and 3 consecutive lost packets
RTPRcvrLostPackets_N4	Number of 4 consecutive packets lost

Table 7. RTP Metrics in the Brooktrout Bfv API

For G.711 on a well managed network, the metrics in Table 7 should report as zero. If non-zero values are being reported, the source of the packet loss should be identified and, if possible, eliminated.

Figure 10 is an example showing how the “RTPRcvrLostPackets_NS” metric behaved under different levels of single packet loss for a short single page fax. Note that as the percent of single packet loss increased, the number of lost packets reported by the NS metric correspondingly increased. Because the non-ECM case sent fewer packets overall, it had less exposure to the impairment and reported fewer dropped packets, as expected.




Figure 10. RTP Packets Lost NS Metric Behavior Under Varying Packet Loss

If jitter is a suspected source of lost packets, the “RTPRcvrJitter” metric indicates the average jitter observed for a G.711 RTP stream. If the reported average jitter value approaches the level of the configured jitter buffer setting, this could indicate a source of lost packets.

Metrics within the Brooktrout Bfv API are also available to use for detecting cases of high latency, such as those shown in Table 8.

Latency Metric	Value Being Measured
time_t4	Maximum time to receive responses to commands, in milliseconds
time_t38	Time to switch into T.38, in ms

Table 8. Latency Metrics in the Brooktrout Bfv API

The “time_t4” metric shows the latency the Brooktrout SR140 experienced during each call. This includes the network latency, as well as the processing time of the remote endpoint to react to a command.

In Figure 11, note that as the round trip latency was increased, the time_t4 value showed a linear increase until it reached its maximum of three seconds. If time_t4 consistently reports values nearing three seconds, this indicates a latency issue in the network that should be investigated.

Figure 11 also illustrates that specific metrics are only applicable to certain phases of a fax call. If a fax fails in an earlier phase of a call than is applicable to a metric, then that metric will typically be zero. In Figure 11, note also that when calls began to fail due to high latency, the time_t4 value was zero. As another example, if there was high packet loss and the call control failed before fax transmission could begin, then the lost packet statistics would be zero because they only apply to the fax portion of calls.




Figure 11. time_t4 Metric Behavior Under Varying Round Trip Latency

The “time_t38” metric is another indicator of network latency. In Figure 12, note that time_t38 increased in direct proportion to the latency on the network, and was able to report values beyond which faxes will fail due to that latency. Values returned by time_t38 are network-specific, because different brands of gateways will have different timing. Figure 12 also shows that time_t38 was only applicable to T.38 mode and not G.711.




Figure 12. time_38 Metric Behavior Under Varying Round Trip Latency

Numerous other fax quality metrics are provided by the Brooktrout Bfv API, which can be useful for diagnosing network issues. Refer to the API reference manual listed in the “For More Information” section for more details.

Appendix B—Test Methodology

Fax Sending Machine

Dell PowerEdge R200, Intel Xeon dual core CPU @ 3 GHz, 4GB of memory, running Windows Server 2003 Enterprise Edition with Service Pack 2.

Fax Receiving Machine

Dell PowerEdge 1950, Intel Xeon quad core CPU @ 3 GHz, 2GB of memory, running Windows Server 2003 Standard Edition with Service Pack 2.

Test Application

FDTool (Dialogic® Brooktrout® Fax Diagnostics Tool) v1.3 with enhancements to calculate and record fax quality metrics, built with Dialogic® Brooktrout® SDK Version 6.4.0.

Configured to use the Dialogic® Brooktrout® SR140 IP Host-based Fax virtual module, with SIP call control, sending a one page fax on five channels for 10 cycles (that is, 50 faxes per data point).

Impairment Tool

NetDisturb Impairment Tool for IP Networks – Enhanced Edition, Version 4.8.0, running on the fax sending machine using Broadcom NetXtreme Gigabit Ethernet cards.

For packet loss tests, the UDP traffic was impaired bi-directionally using the “Loss: Percentage” mode of the Tool, which uses a random packet loss distribution algorithm.

For burst loss tests, the UDP traffic was impaired bi-directionally using the “Loss: Percentage & Duration” mode of the Tool, configured for one second intervals with 100% loss during the burst loss duration.

For latency tests, the UDP traffic was impaired bi-directionally using the “Constant Delay” mode of the Tool.

References

[ITU-T] ITU-T Recommendation G.1050 - Network model for evaluating multimedia transmission performance over Internet Protocol

For More Information

Dialogic® Brooktrout® Bfv APIs Reference Manual

ITU-T Recommendation T.30 - Procedures for document facsimile transmission in the general switched telephone network

ITU-T Recommendation T.38 - Procedures for real-time Group 3 facsimile communication over IP networks

ITU-T Recommendation E.450 - Facsimile quality of service on public networks – General aspects



www.dialogic.com

Dialogic Inc.
926 Rock Avenue
San Jose, California 95131
USA

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH PRODUCTS OF DIALOGIC INC. AND ITS AFFILIATES OR SUBSIDIARIES ("DIALOGIC"). NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN A SIGNED AGREEMENT BETWEEN YOU AND DIALOGIC, DIALOGIC ASSUMES NO LIABILITY WHATSOEVER, AND DIALOGIC DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF DIALOGIC PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHT OF A THIRD PARTY.

Dialogic products are not intended for use in certain safety-affecting situations. Please see <http://www.dialogic.com/about/legal.htm> for more details.

Dialogic may make changes to specifications, product descriptions, and plans at any time, without notice.

Dialogic and Brooktrout are registered trademarks of Dialogic Inc. and its affiliates or subsidiaries. Dialogic's trademarks may be used publicly only with permission from Dialogic. Such permission may only be granted by Dialogic's legal department at 926 Rock Avenue, San Jose, California 95131 USA. Any authorized use of Dialogic's trademarks will be subject to full respect of the trademark guidelines published by Dialogic from time to time and any use of Dialogic's trademarks requires proper acknowledgement.

The names of actual companies and products mentioned herein are the trademarks of their respective owners. Dialogic encourages all users of its products to procure all necessary intellectual property licenses required to implement their concepts or applications, which licenses may vary from country to country.

Any use case(s) shown and/or described herein represent one or more examples of the various ways, scenarios or environments in which Dialogic® products can be used. Such use case(s) are non-limiting and do not represent recommendations of Dialogic as to whether or how to use Dialogic products.