

Kofax Capture 9.0 Citrix and Terminal Services Considerations

(Kofax Capture 9.0 Enterprise with Citrix & Terminal Services FAQs)

Application Note

Date	July 18, 2010
Applies To	Kofax Capture Enterprise (Only) 9.0, supported with: <ul style="list-style-type: none">• Terminal Services with Windows Server 2003 Enterprise• Citrix MetaFrame Presentation Server 3.0• Citrix Presentation Server 4.0, Standard• Citrix Presentation Server 4.5, Enterprise• Citrix XenApp 5.0
Summary	This application note provides answers to frequently asked questions when implementing Kofax Capture 9.0 Enterprise on Citrix & Terminal Services.
Revision	3.0

Kofax Capture Enterprise Deployment Considerations

The purpose of this document is to provide guidance when consideration is being given to running Kofax Capture Enterprise version 9.0 on Citrix or Terminal Services. As no two infrastructure environments are the same, each environment must be evaluated separately, based on systems in use, network speed and type, the functionality planned or being used within Kofax Capture, the response time expected by end users, etc. Due to each of the aforementioned variables that contribute to corporate infrastructures, we are responding to the most frequently asked questions and technical areas that should be evaluated and analyzed when deploying Citrix or Terminal Services with Kofax Capture Enterprise 9.0.

Frequently Asked Support Questions

1. What is the purpose of Citrix and Terminal Services?

Citrix and Terminal Services technologies allow you to run software (such as Kofax Capture) on a centralized server (or bank of servers), and have the user interface (UI) rendered on clients with little or no software installation.

These benefits apply mainly to larger, more complex sites:

- IT overhead is reduced because software installation and deployment occurs on a single server (or a small bank of servers) in a centralized location. There is no need to remotely deploy or update software (like Kofax Capture).
- Speed of deployment. If deploying to a number of machines (especially if they are remote), then these technologies allow the users to gain access to Kofax Capture Enterprise 9.0 almost immediately.
- Security is also increased because the client machines do not need software installation and can be left empty and “locked down.”

2. Can I use this feature without Kofax Capture Enterprise?

With Kofax Capture, you are limited to one login session with Citrix or Terminal Services. This permits some basic testing and administration, but is not feasible for production usage. Kofax Capture Enterprise 9.0 does not limit login sessions.

3. How does Licensing Work?

When connecting with these technologies, each login, or “session,” acts like a separate station, and thus requires an appropriate station license (such as a Full Station). As noted above, if you intend to connect to more than one session, then you will need Kofax Capture Enterprise 9.0.

4. Should I choose Terminal Services or Citrix? What is the difference?

Terminal Services is included with Windows Server 2003 Enterprise. Citrix is essentially an upgraded version of the same feature with additional capabilities, including server pools, a Web-based interface, and additional administration tools. Citrix is separate product with additional cost. Kofax does not recommend the use of one over the other. This is dependent upon the use requirements of the organization and budgetary considerations.

5. What are the supported product versions?

Kofax Capture Enterprise 9.0 is supported by Kofax with:

- Citrix MetaFrame Presentation Server 3.0
- Citrix Presentation Server 4.0 Standard Edition
- Citrix Presentation Server 4.5 Enterprise Edition
- Citrix XenApp 5.0

6. Which Kofax Capture Enterprise modules are supported?

Starting with Kofax Capture Enterprise 9.0, all built-in Kofax Capture modules are certified. Custom Modules and add-on products may require changes to support this environment. Please contact the vendor of your Custom Module, if applicable, for compatibility.

7. Is scanning supported?

No. The application is actually running on the server, so you cannot use a client-connected scanner.

8. What sort of network connection do I need?

Terminal Services and Citrix work by sending all screen updates to the client computer. On applications such as Microsoft Word or Excel, the level of screen updates are small in a given time period, and adequate performance can be achieved on a very low-speed network.

On the other hand, an image-intensive application like Kofax Capture Enterprise 9.0 needs to update large parts of the screen when displaying images. This is compounded with the typical need for very fast key entry at Validation Stations.

Because of these requirements, deploying an imaging-based application like Kofax Capture with Citrix or Terminal Services requires a network that can accommodate a lot of data at once (high bandwidth), as well as minimal latency. A lower average rate of response time in terms of Internet traffic is also required to prevent performance related issues, especially on offshore Citrix clients. In terms of packet loss, make an assessment of your network and ensure that packet loss is 5% or lower. Otherwise, connections may become slower and latency will increase due to multiple packet retries (attempts to resend the information).

9. Can users authenticate using Active Directory?

Yes, users can authenticate using Active Directory. However, authentication using Active Directory may cause Citrix-based users to notice a substantial delay. Testing should include users from all Active Directory forests and/or trees. This will ensure that all Citrix users validating from different Active Directory servers have an acceptable level of performance.

10. How much network bandwidth do I need?

Bandwidth is a measure of the amount of data that can be sent across the network in a given period of time.

The Kofax Capture applications most likely to be used across Citrix are Quality Control (typically to review images), and Validation and Verification (to do batch editing and/or data-entry).

These applications involve combinations of the following operations:

- Typing in a set of keystrokes for a field, and tabbing to the next field or document. This typically produces a screen update that shows a zoomed version of the next field.
- Clicking or utilizing a keystroke on some element of the batch contents tree to perform a batch editing operation or review a page. In either case, the screen typically updates to show a large version of the selected image.

Both of these operations are similar in that they initiate an image update, which may cover a large portion of the application window.

When using Citrix, such an image update means that a compressed version of the entire displayed image (bitmap) must be sent across the network. This bitmap update, which is common to all imaging applications, is the crux of frequent configuration mistakes with Citrix because it is so different than Microsoft Office and similar desktop applications.

Another common mistake is to take into account only single-user performance when calculating bandwidth. Because multi-user tests are expensive and time consuming, some users only test with a single user or a small set of users. When the full complement of users is brought on line, the combined volume of the bitmap refreshes can overload the network.

The following procedure suggests a way to plan likely bandwidth requirements for Kofax Capture Enterprise 9.0 without the necessity of performing multi-user tests.

NOTE: Performance tests and sample data are provided based on Citrix 4.5 Enterprise.

Screen Refresh Estimate

The first step is to estimate the amount of data transferred for a bitmap refresh. Using Kofax Capture Validation (maximized) with the default layout on a 1280 x1024 screen, Kofax has found the following averages for the number of bytes transferred to a Citrix client for the given operations:

Magnified Zonal Image (such as shown for an index field)	Black & White	30 KBytes
	Color / Gray	1 MByte
Full page display (such as when doing batch editing or performing image quality control)	Black & White	100 KBytes
	Color / Gray	400 KBytes

The above numbers are related to using images with forms and/or significant textual content. When using less dense images, the amount of data sent can be as little as half.

Citrix appears to compress the data sent based on the variance of patterns as well as the density and variance of color content.

For black & white images, full page displays have a relatively high amount of variance (i.e., showing the entire form), while magnified images display less variance (because only a small portion of the form is shown).

In the case of color images, a full page display tends to combine much of the color information (especially if much of the form is black & white text). A magnified color image will tend to show color variation much more substantially, and thus it requires much more bandwidth.

If desired, you can set up a test and run Kofax Capture Validation (or Quality Control) through Citrix with the aid of a network sniffer. In Validation, switch fields (or click on a new page or document), and identify the

number of bytes sent in the utility. Ideally, this test would be customized to your specific application scenario.

NOTE: It is beyond the scope of this document to identify specific utilities or provide detailed testing procedures. If you do not feel comfortable performing this test, then you should consider hiring an outside resource experienced with Citrix and networking.

The subsequent examples utilize 30 KBytes as a reasonable number of transferred bytes for validation of fields on a form.

Operations per Minute

The next step is to estimate the number of operations executed per second. It is suggested that the average time to key a field is somewhere between 1-3 seconds (for a high speed operator). Image review time may be faster (from .5-1 second). If additional data in this area is unavailable, it is suggested to use 1 operation per second as a rough estimate.

If you have a high speed operator, then you can have that operator process several test batches (without Citrix) with your documents.

For a quality control estimate, take the number of images reviewed (i.e., displayed) divided by the time to complete the test.

For a validation estimate, take the number of fields on each document times the number of documents divided by the time to complete the test. Recall that the primary bandwidth usage is from the magnified zone refresh, not the key entry. Thus, it is the number of fields (i.e. the number of zone refreshes) that is critical to count, not the number of characters typed. A point to consider is whether some fields will be automatically skipped due to OCR validation, scripting, or other means. In such cases, you need to estimate the average number of fields that will actually be displayed (rather than all fields).

Number of Operators

You must now consider the number of operators expected at peak usage (and allow for some likely growth).

Other Network Usage

You must also add in other network usage. Be sure to measure this at peak times and consider other applications that may be coming on line in the future.

Final Computation

A final computation of bandwidth is as follows:

$$((\text{Screen Refresh Data Size}) * (\text{Operations per Second}) * (\text{Number of Operators})) + (\text{Other Network Usage})$$

An example using default measurements with 10 operators (with no other network usage) might be:

$$0.03 \text{ MBytes} * 1 \text{ ops/sec} * 10 \text{ operators} = 0.3 \text{ MBytes/sec} = 2.4 \text{ Mbites/sec}$$

Is it Enough?

Once you have computed your likely bandwidth, you need to compare it to the bandwidth rating on your network. Consult your network administrator for this measurement.

If your required bandwidth isn't sufficient, then you will need to scale up your network infrastructure appropriately to avoid potential performance degradation. If the network is not scaled up, then it will essentially limit the amount of data sent, which will negatively impact the speed of processing.

11. Can I tune Kofax Capture Enterprise to reduce bandwidth usage?

Yes. In the equation presented above, the key variable that can be altered is the screen refresh estimate. The number of bytes required for such a screen refresh can be reduced through the following techniques.

Reduce the Kofax Capture Window Size

The simplest approach is to reduce the Kofax Capture window size, which will reduce the amount of data sent because the image updated is smaller.

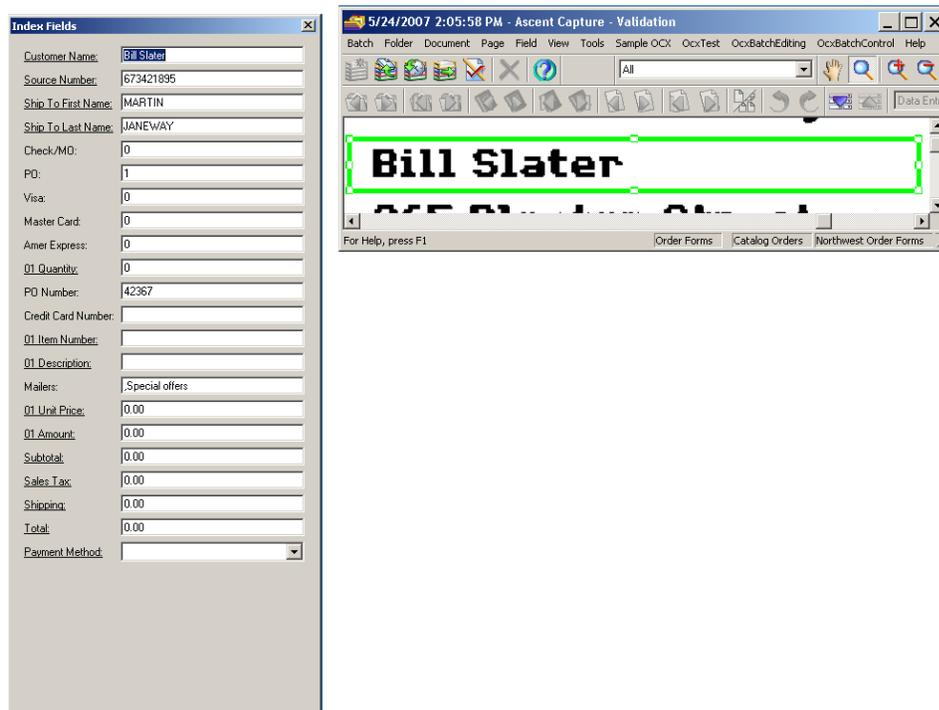
In tests at Kofax, reducing a maximized window (1280x1024) down to approximately 500x300 decreased the bytes sent per operation to only 30% that of the full-screen version. The average operation required roughly 10 KBytes (instead of 30 KBytes with full screen)

Reduce the Image Viewer Size

Reducing the overall window size may not be ideal because users might need to see a larger view of the batch tree or the data entry panel (both of which don't change much and don't have a large impact on bandwidth).

To address this, you can undock the batch contents panel or the data entry panel, size them appropriately, and then minimize the main window (which typically contains the image viewer). The Scan Controls panel (in Quality Control) can also be undocked. For instructions on docking and undocking, see the Kofax Capture Help.

An example layout might look as follows:



The reduction in the number of transferred bytes with this option will be similar to when you simply reduce the overall window size.

Turn off automatic zoom and pan in Validation

Normal operation in Validation and Verification is to pan and zoom the image view so that a magnified view of the current zone is displayed at the upper-left. This tends to cause the entire image viewer to refresh (and thus a large amount of bandwidth is used with Citrix).

An option in some cases may be to enable full page view on the viewer and disable the pan and zoom behavior. After doing so, the system will show different highlighted areas on the screen, but will not repaint the entire image viewer except when moving to different documents and pages.

In such cases, it must be possible for the operator to easily read all the data fields without the need to zoom. Operators have the ability to manually pan as well, but this is strongly discouraged under Citrix since panning will generate significant network bandwidth usage.

Tests with a full screen (1280x1024) validation application with either a color or black & white image yielded operations of about 15 KBytes (roughly 1/2 the normal amount for black & white images, and a significant savings for color images).

A caveat is that each new page displayed (typically one or more per document) will require a full refresh which, as indicated above, takes about 100 KBytes/image.

Thus, this technique might only be advisable when there are a large number of index fields per page.

12. How low does my network latency need to be?

Latency refers to the amount of time it takes to send a particular packet across the network. Low latency is good, while high latency is bad.

There are two likely places with which to be concerned:

Key Entry

There is some interaction between the client and server between the time a key is pressed, and the associated character is displayed on the screen. If latency is too slow, then there will be a perceptible delay.

If your network latency is high, this issue will show up with applications such as Microsoft Word. A simple test is to launch Word from a Citrix client (at a remote location), instruct a fast typist to enter a paragraph, and report if any delays are noticed while typing.

If delays are noticed, then this should be considered a severe problem with the network. Keying from image applications like Kofax Capture tend to be data-entry intensive, and high latency will impact productivity.

Screen Refresh

Another area that intersects latency is screen refresh. Typically, a user in Kofax Capture clicks on a page or document, or tabs between fields, which causes a refresh of one or more page images. Such operations occur less frequently than keystrokes, and thus latency has a smaller impact. Generally, if the key entry test above passes acceptably, then screen refresh should not be impacted by latency. If screen refreshes are slow, then bandwidth is likely the culprit (see above).

General Latency Testing and Guidelines

We recommend that your network staff provide latency testing with the network tools they normally use. However, a rough test can be performed with the Ping command.

To do this, do the following from the remote client:

1. Select the Windows "Start" Menu.
2. Select "Run..."
3. Type "cmd" and click OK.

4. This pops up a command window.
5. Type “ping (citrix server name)” – An example might be “ping www.google.com”

An example display might be similar to the following:

```
Pinging www.1.google.com [72.14.253.104] with 32 bytes of data:  
Reply from 72.14.253.104: bytes=32 time=38ms TTL=237  
Reply from 72.14.253.104: bytes=32 time=36ms TTL=237
```

After a few seconds, hold down the “Ctrl” key and the “C” key at the same time to stop the command.

This command sends a simple network packet to the server and back and times the response, which is essentially a measure of the latency.

In the example above, the round trip time (RTT) is 36-38ms. Based on customer experience, round trip latency above 200ms has been seen to impact productivity, and Kofax recommends necessary infrastructure to keep latency below this threshold.

NOTE: Your administrator may have disabled the Ping command on the server, in which case, this technique cannot be used.

SpeedScreen

Citrix 4.5 includes a new feature called SpeedScreen which optimizes display.

Bandwidth tests with Kofax Capture screen refreshes did not show any significant difference in the amount of data transferred with SpeedScreen enabled or not. However, these tests were explicitly designed to avoid repeated display of images and/or scrolling. In a real usage scenario, it is possible that SpeedScreen may provide some level of transfer savings. Kofax testing did not find any issues with this feature, so it is recommended to leave it on.

Publish Applications (Application Virtualization)

Citrix has two primary modes for displaying applications:

Remote Session – This displays an entire desktop where users can launch various applications.

Application Virtualization – The user launches a specific application, and only the application window is displayed on the remote client.

Kofax tests both modes, but strongly recommends Application Virtualization. This mode requires fewer server resources (since only the application is launched, not the entire desktop), and requires less network bandwidth (since the system does not need to transmit refreshes to the desktop).

13. What Citrix Features are known to be incompatible with Kofax Capture Enterprise?

These features are known to be incompatible with Kofax Capture Enterprise 9.0 at this time:

Virtual Memory Optimization (“Rebasing”)

This feature does not currently work with Kofax Capture Enterprise 9.0. We are considering further research in consideration of the priority of other features.

256 Color Mode

Kofax Capture Enterprise 9.0 does not support 256 color mode due to various graphical display requirements. There are no plans to offer such support in the future.

Some users want to utilize this mode because they believe that it saves on bandwidth.

Tests by Kofax with Kofax Capture Enterprise 9.0 have shown that 256 color mode over Citrix does not reduce bandwidth requirements. Based on these results, it is surmised that Citrix optimizes and compresses the screen updates so that if you are displaying a black & white image, it only sends black & white, no matter what the screen depth is.

If it is for some reason necessary to run other applications in 256 color mode, it is suggested to separately publish Capture applications using application virtualization. In this way, separate applications can have different color settings.

14. What alternatives can I use if I have insufficient bandwidth, or my latency is too high with Citrix?

If you are not able to scale your network to appropriately support the number of desired users of Kofax Capture Enterprise 9.0 on Citrix, it is suggested to deploy KCN Server. Although installation of an application is required at the remote site, KCN Server can function with extremely high performance over slow connections since it uploads/downloads data asynchronously to the interactive application.

15. Should performance of non-interactive operations be identical under Citrix as compared to a standard workstation?

Yes. Citrix should only impact operations that relate to interactive operations such as displaying images, or keying data. Other operations such as opening and closing batches, running scripts, and validating against databases should run at an identical performance level as compared to running on standard workstations.

If this is not the case, then there is most likely an issue with the configuration of the Citrix server itself related to memory limits, CPU bandwidth, or network infrastructure to the Capture server or some other critical resource.

16. How many instances of Validation, Verification, or QC can I run on a single Server?

The number of instances per server can vary substantially based on CPU, Memory, network setup, and actual use-cases, and other factors.

Successful customer use-cases are known with from 5 to 25 instances per server.

As a rough guideline, Kofax set up tests on a single CPU, 3.0 Mhz (HyperThreaded) server with 2GB of RAM. From 1 to 30 Citrix sessions were started on this machine, and proceeded to process batches consisting of 360 bitonal pages, 20 documents, with 5 index zones per document. Each session was run interactively with a script that simulated acceptance of a file after ½ of one second.

The results show that degradation is relatively minor (less than 10%) even at 30 users.

1 session:	1.42 seconds per batch
2 sessions:	1.42 seconds per batch
5 sessions:	1.42 seconds per batch
10 sessions:	1:44 seconds per batch
20 sessions:	1:48 seconds per batch
30 sessions:	1:53 seconds per batch

With 30 users, the CPU usage never jumped above roughly 75%, and swapping activity was minimal. Processing time for each batch was measured from batch open to batch close time.

Using a more powerful server, for instance with multiple CPU's and more memory, should allow for more instances (or for less degradation) than seen here.

Different operations (such as using CPU-intensive scripts or substantial database lookups) may reduce the number of instances that can be supported on a machine.

17. What are the network requirements between the Citrix server and the Kofax Capture server

The Kofax Capture Server must be reachable over a LAN (Local Area Network) from the Citrix Server. In cases where Kofax Capture Enterprise 9.0 components are distributed to different machines (such as alternate database servers and image file servers), all such machines must also be reachable via the LAN from the Citrix Server.

Kofax Capture Enterprise 9.0 does not support deployment of a workstation client or a Citrix server over a WAN (Wide Area Network) as compared to the Kofax Capture Server.

18. Do I need to purchase additional licensing to install Report Viewer in a Citrix environment?

No additional license is needed for Kofax Capture but a full licensed version of Crystal Reports XI (version 11.0.0.0895) is required to be installed in the same machine. That is due to a Crystal Reports licensing requirements.