


The Pathomation software platform for digital microcopy version 2: system requirements

© 2019 Pathomation

The Pathomation software platform for digital microcopy version 2: system requirements

by Angelos Pappas, Yves Sucaet



The Pathomation software platform for digital microcopy version 2: system requirements

Copyright 2019, Pathomation

All rights reserved. No parts of this work may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without the written permission of the publisher.

Products that are referred to in this document may be either trademarks and/or registered trademarks of the respective owners. The publisher and the author make no claim to these trademarks.

While every precaution has been taken in the preparation of this document, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of information contained in this document or from the use of programs and source code that may accompany it. In no event shall the publisher and the author be liable for any loss of profit or any other commercial damage caused or alleged to have been caused directly or indirectly by this document.

Printed: December 2019 in Berchem, Belgium

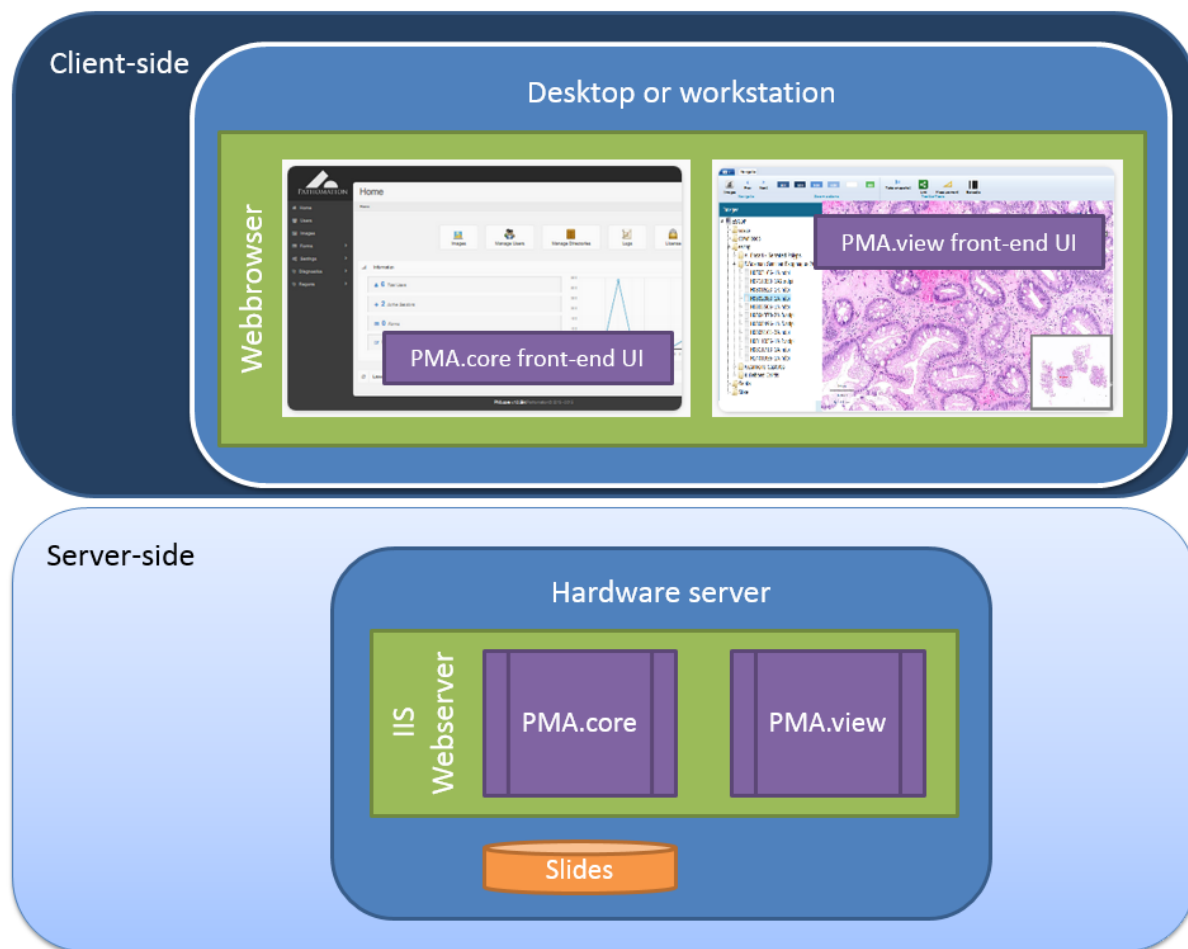
Table of Contents

Foreword	0
Part I Introduction	5
1 A platform.....	5
Part II Scenarios	6
1 Single server.....	6
2 Two servers.....	7
3 Supported storage media.....	7
Part III System Requirements	9
1 Client-side	9
Desktop or workstation hardware	9
Web browser	9
Popup-blocking and Ad-blocking.....	10
Supported hosting applications	11
2 Server-side	11
Server hardware	11
Operating system and software prerequisites	12
3 Configuration hints and scaling considerations.....	12
4 Notable changes between version 1 and version 2 of our platform.....	13
Part IV Performance	14
1 Network transfer.....	14
2 File format comparison.....	14
Index	0

2 Scenarios

Below we illustrate two different scenarios to deploy our software. The details of these will vary depending on your specific needs, but can still be distilled from the information given herein.

2.1 Single server

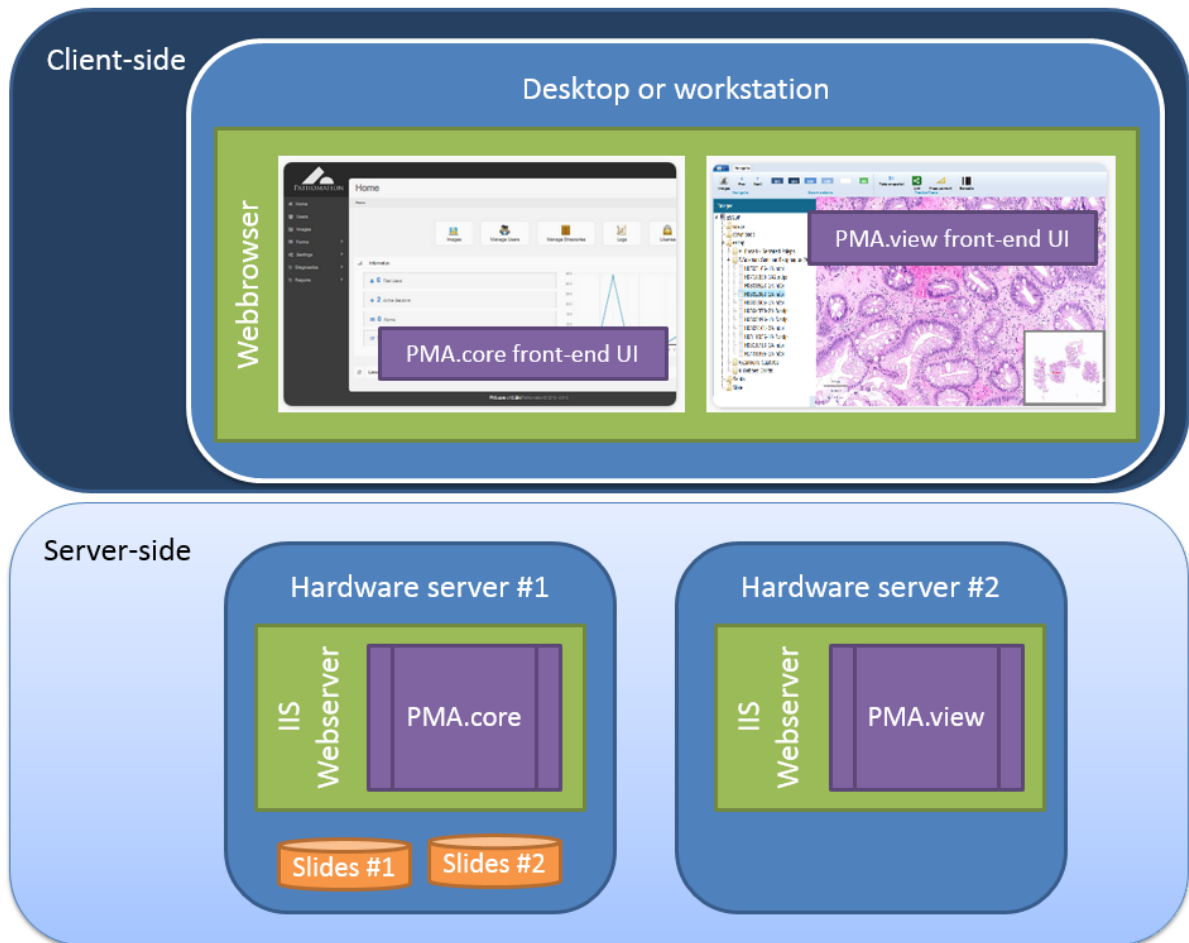


A single piece of hardware is used as a server to run the IIS webserver software, which in turn hosts both PMA.core and the downstream viewing application PMA.view. A single slide repository is available on this machine as well.

Both PMA.core and PMA.view have a front-end user interface (UI). This front-end can be accessed via any computer in the organization's network, or even via the Web (possibly after establishing a VPN to the organization's internal LAN first).

As PMA.core and PMA.view run on the same server hardware, requirements can be summed up as "the maximum of the two". The hosting of PMA.core requires more resources than PMA.view, but when both application are placed on the same hardware, PMA.core requirements should suffice to host both.

2.2 Two servers



Two pieces of server hardware are available in this setup and both run the IIS webserver software. The server that hosts PMA.core also serves as a file server with two slide repositories.

Both PMA.core and PMA.view have a front-end user interface (UI). This front-end can be accessed via any computer in the organization's network, or even via the Web (possibly after establishing a VPN to the organization's internal LAN first). The only difference with the previous setup is that PMA.core will be accessible through <http://server1/...>, while PMA.view will be accessible through <http://server2/...>

As PMA.core and PMA.view run on the different server hardware, they can have different resources. The hosting of PMA.core requires more resources than PMA.view, so server #2 need not have as many resources as server #1.

2.3 Supported storage media

PMA.core can access whole slide images located in any type of storage compatible with Microsoft Windows, whether that is a regular disk drive, a network share or a flash memory stick. On top of that, PMA.core is capable of accessing whole slide images stored in Amazon S3 compatible storage. Each of these options comes with it's own advantages and disadvantages.

Locally attached disk drives

This is probably the most common scenario. Whole slide images can be stored in a disk drive which is physically attached to the machine hosting PMA.core. The type of the drive greatly affects performance, but can also increase storage costs or raise durability issues.

- **SATA** drives are the most common. They are mechanical drives, resembling a pickup. They are relatively cheap, have high durability, moderate performance and come with large storage capacity.
- **Solid State Drives** are a more recent technology. They are a mixture of RAM and flash drives, offering superior performance. The downside is the lower durability, higher cost and smaller capacity compared to SATA drives.

Network shares / drives

PMA.core can mount network shares and also access slides coming from network drives. This option gives greater flexibility in terms of storage as it allows the addition of storage that doesn't have to be physically attached to the host machine. In such setups the performance depends on the connectivity between the two parties, as well as the type of the networked storage hardware. Usually Gigabit Ethernet or faster is suggested.

Amazon S3 compatible storage

Amazon Simple Storage Service is a web service storage where data access happens over HTTP requests. The benefits of using Amazon S3 as a storage are the inherent capacity scalability and cost when storing large volumes of data. The downside is that its performance is usually worse compared to a locally attached drive.

PMA.core cache requirements

Regardless of the choice or combination of the storage media used, the system requires some space on a local or network disk where cache data will be stored and accessed. PMA.core allows you to define the cache retention policy and informs you about its size. As a general rule of thumb, the cache capacity should be around 10% of the total size of the available whole slide images.

3 System Requirements

The platform consists of a web application and an optional system service for cache scheduling. The software is built for the Microsoft Windows platform and is capable of running in both user and server versions of Windows.

The user interface for our software runs in a webbrowser. Clients are therefore expected to be able to interact with our system in almost any major modern browser with JavaScript enabled on any modern Mac, PC (Windows or Linux). We offer only limited support on mobile platforms for the moment, and features may only be expected to work to the extent that the mobile browser being used is a successful port of its own desktop version.

Requirements and recommendations are given under the assumption that a machine is used for a single purpose only. If you're using a hardware server as a hosting environment for PMA.core as well as a file server for your digital slides, your network bandwidth requirements e.g. will be significantly higher than if you only used it for hosting PMA.core.

Please be aware that system load and performance depend highly on your particular use case for digital pathology. There is a significant difference between hosting a static collection for a high number of concurrent users (typical in an educational setting) or hosting a high-turnover collection that is only consulted by a limited number of pathologists. As it is, we make several recommendations on how you can tune your setup for your particular working conditions.

3.1 Client-side

During development our software is continuously tested on the Microsoft Windows Operating System using the following browsers: Internet Explorer 10+ and the latest versions of Google chrome, Safari and FireFox.

3.1.1 Desktop or workstation hardware

The minimum requirements in terms of processor speed and memory (for end-users) are hard to determine. We are pretty sure that every computer younger than 5 years will be able to interface with the platform using a webbrowser without any problems. In case of doubt, we refer to the particular system requirements that apply to the specific version of flavor of webbrowser that is being used.

3.1.2 Web browser

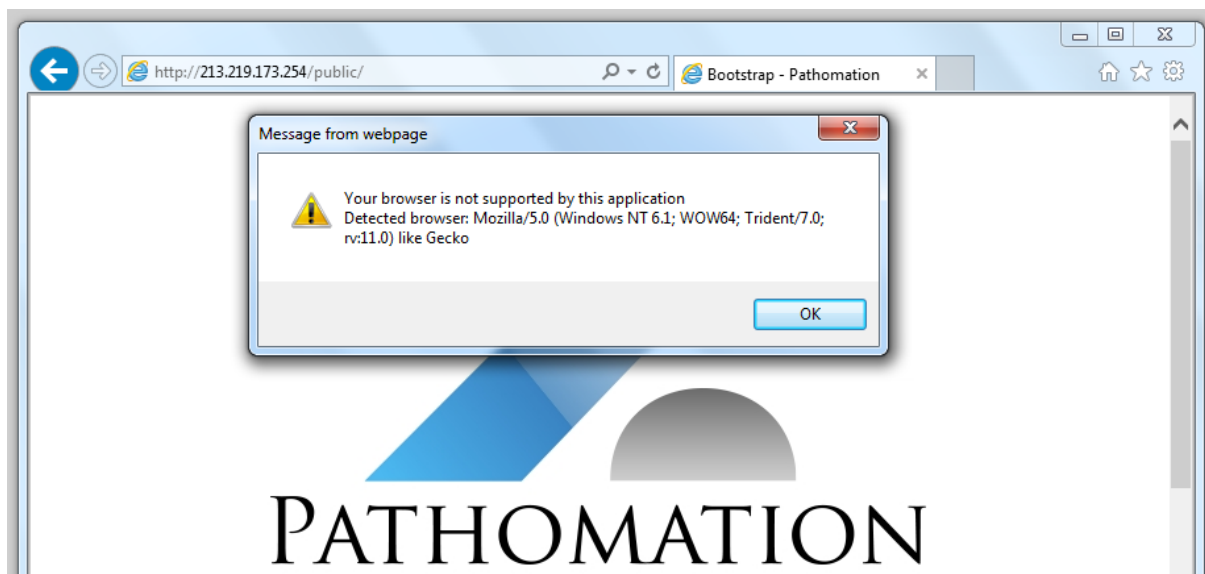
A hosted web application is consulted through a (usually) remote computer on the internal company network (or Internet). To allow the delivery of content to as many clients as possible, we have very relaxed requirements in terms of client hardware and software requirements.

Key is that we don't require any additional webbrowser-plugins such as Flash, Silverlight, WebGL, or Java. Nor do we depend on or do we install any additional browser add-ons (BHO - Browser Helper Objects) ourselves.

You will need a modern HTML5-capable webbrowser. This means that your browser must be configured to allow JavaScript and accept cookies. Besides that, we've tested our software extensively on a host of different web browser environments, and found it to operate under almost all, except Internet Explorer 10 or lower. For the best user experience, we recommend using Mozilla Firefox, Microsoft Edge or any

Chromium based browser such as Google Chrome and Opera. On a Mac, you are welcome to use Safari, Firefox or, again, any Chromium based browser.

When a browser is not supported, an error message follows.



Because of content delivery through a web browser, interfacing with our software is not limited to Windows-only machines (though the web applications themselves run on Microsoft infrastructure). Our communication implementation is based on REST web services and the standard HTTP protocol, so you can truly use our software from anywhere, any time.

3.1.2.1 Popup-blocking and Ad-blocking

Our software uses dialogs as a means to interact with the end-user. These dialogs are implemented based on "best practice" industry standards (including libraries such as React and jQuery). These dialogs are not pop-ups, and during testing, no interference with known pop-up blockers was ever observed.

That being said, it is possible that at some point in the future these pop-up blockers pick up features in our site and label them as popups, subsequently rendering them ineffective. If this happens, you should let us know and we will help resolve the issue in collaboration with you popup blocking software vendor.

The same reasoning applies to webbrowser plugins that block advertisements. These oftentimes already target the kind of technology that we use ourselves to interact through dialogs (such as iframes). There is a large variety of possible interpretations of such features (in a website), and while we haven't come across this yet, it is possible that at some point a certain feature of our site gets interpreted as commercial content. If this happens, you should let us know and we will help resolve the issue in collaboration with you advertisement blocking software vendor.

If you experience problems on your installation, please let us know and we will find a solution to resolve your particular issues.

3.1.3 Supported hosting applications

Our platform offers optional plugins for various host applications such as Adobe Photoshop and ImageJ. The system requirements for these external applications are to be specified by their respective vendors. We don't place additional demands on top of the standard manufacturer's recommendations.

As a rule of thumb, when our PMA.view web-based viewer client is performing well on a user's computer, it can be expected that the plugin for the host application under consideration will perform equally well.

3.2 Server-side

Note that the platform consists of the PMA.core (server-side) web application and the PMA.view (end-user; client-side) viewer web application. Both can be installed side by side as different web applications on the same server. We actually recommend this setup initially, since the viewer web application was designed to have a very small footprint (thin server / fat client concept).

3.2.1 Server hardware

For an initial setup of PMA.core (adequate to serve most conditions) we recommend the following:

Component	Minimum	Suggested
CPU	2.2 GHz Quad Core or AMD equivalent	Intel® Core™ i5-4200H 3.40 GHz or AMD equivalent
Number of CPU cores	4	8
RAM	8GB DDR3	16GB DDR3
Hard Disk	SATA II	SATA III SSD for Operating System
Network bandwidth	100 Megabit	1 Gigabit

In a subsequent phase, PMA.view may be installed on a separate machine (see [Two server scenario](#)). As the heavy work is done by PMA.core, the specific requirements are significantly less:

Component	Minimum	Recommended
CPU	Intel i3 3220 or AMD FX-4300	Intel i7-47xx or AMD FX
Number of CPU cores	2	4
RAM	8GB	16GB
Hard Disk	SATA II	SATA III
Network bandwidth	100 Megabit	1 Gigabit

3.2.2 Operating system and software prerequisites

In principle any version of Windows + Internet Information Server (IIS) will do. We've got the software to run on IIS Express edition for demonstration purposes. We offer the following formal guidelines:

Component	Minimum	Suggested
Operating System	Windows 2012 Server	Windows 2016 Server
Web Server	IIS 7.5	IIS 10.0
.Net framework	4.5	4.7
RDBMS*	Microsoft SQL Server 2014 Express	Microsoft SQL Server 2017 Express

* SQL Server can be deployed in a separate machine or even a cloud service such as Windows Azure or Amazon RDS.

If you want either software component of the platform to interact with end-users through email, you will also need to provide an SMTP-server. This server does not need to be installed on the same machine as the Pathomation software.

3.3 Configuration hints and scaling considerations

Improving disk I/O throughput

The system's response time can be greatly improved by installing it on a solid state disk. The cache directory should also be kept on an SSD as well. Thus an optimal setup would be to install the operating system and the application on an SSD while storing digital slides in regular hard disks (or network shares). Do not underestimate the importance of the regular disks' performance, where slides are stored, though. Whenever possible use 6Gbit SATA III disks.

Network

Multiple NICs may be installed onto the server machine to increase network bandwidth. The system can utilize multiple IP addresses or host names out of the box, thus special load balancing configurations are not required.

Parallel processing

The system is built to process requests in parallel and its performance depends on the available cores per processor, the performance per core, as well the total number of processors in the system. In environments with more than 50 simultaneous users, two processors should at least be available in the system.

3.4 Notable changes between version 1 and version 2 of our platform

1. Switched to Kakadu for JPEG 2000 rendering
2. Switched to SQL Server as the primary RDBMS
3. Added support for Amazon S3 storage
4. Improved rendering performance
5. Added support for fluorescent MRXS slides
6. Added support for fluorescent CZI slides
7. Added support for fluorescent NDPIS slides
8. Added support for fluorescent AFI slides
9. Added support for OME-TIFF slides
10. Added support for DICOM slides
11. Added support for Menarini GXP slides
12. Added support for Menarini RAW slides
13. Added support for ObjectivImaging slides
14. Added support for Olympus WebView slides
15. Added support for PerkinElmer QPTiff slides
16. Added support for Sakura slides
17. Added support for SZI slides
18. Improved support for Ventana BIF slides

4 Performance

We have profiled PMA.core for a variety of conditions

4.1 Network transfer

We have profiled the PMA.core to examine its performance characteristics in a networked environment.

The tests were done using the following hardware

- Server: Intel i5 @ 3GHz, 8GB RAM, 128GB SSD, 1TB SATA III
- Gigabit Ethernet
- 802.11g (64MBps) WiFi - a relatively cheap WiFi router bought off the shelf

Measurements using tile quality 75 and tile size 512x512:

Network type	Tile source	Cache building	Performance	Limit source	Concurrent users*
Gigabit Ethernet	Cache on SSD	No (already built)	1000 tiles/sec	CPU	100
Gigabit Ethernet	Native slides on SATA disk	No (bypass cache)	300 tiles/sec	CPU	30
Gigabit Ethernet	Native slides on SSD disk	No (bypass cache)	400 tiles/sec	CPU	40
Gigabit Ethernet	Native slides on SATA disk	Yes (cache build on SSD)	200 tiles/sec	Disk	20
Gigabit Ethernet	Native slides on SSD disk	Yes (cache build on SSD)	300 tiles/sec	Disk	30
WiFi 54Mbps	Cache on SSD	No (already built)	150 tiles/sec	Network	15
WiFi 54Mbps	Native slides on SATA disk	No (bypass cache)	150 tiles/sec	Network	15
WiFi 54Mbps	Native slides on SSD disk	No (bypass cache)	150 tiles/sec	Network	15
WiFi 54Mbps	Native slides on SATA disk	Yes (cache build on SSD)	150 tiles/sec	Network	15
WiFi 54Mbps	Native slides on SSD disk	Yes (cache build on SSD)	150 tiles/sec	Network	15

* Concurrent users is an estimation based on the following assumption: Given a slide of 100,000 x 210,000 pixels (x20 scan) we assume that a user will view 3% of it in period of 5 minutes. Such an image has about 100,000 tiles, thus the user will request 3,000 tiles (3%) in 5 minutes or 300 seconds. This translates to 10 tiles per second per user on average.

Dropping tile quality to 50 improved performance over WiFi to 280 tiles per second.

4.2 File format comparison

Test method

An in-house built profiler tool has been used to test the performance of the various whole slide image formats supported by PMA.core. The measurements were performed on each sample slide separately,

without any cache being available and without creating additional cache during the tests (i.e. cache was totally disabled).

The used sample slides were not always exactly the same across different formats, but similar in dimensions at the extent possible. All of them were encoded in JPEG. 10 simultaneous requests were used for each image.

Server configuration

Hardware & Software

CPU	Intel i5-3330 @ 3.00GHz
RAM	16GB
HDD	SATA III @ 7200 RPM
Network	Gigabit Ethernet
Operating System	Windows 10 Pro
Web Server	IIS 10
PMA.core	Release x64 version

Results

Format	Parser	Dimensions (pixels)	Test run time (hh:mm:ss)	Bytes downloaded	Tiles downloaded	Tiles / sec	AVG tile fetch time (sec)
BIF	PMA.Lib	105948 x 94154	00:02:24	501.75 mb	11776	81.77	0.112
CZI	ZeissImageLib	86609 x 83095	00:01:21	72.14 mb	1369	16.7	0.568

MRXS	PMA.Lib	68608 x 95232	00:01:30	284.53 mb	15663	172.86	0.048
NDPI	PMA.Lib	122880 x 110592	00:01:26	502.44 mb	13653	158	0.053
SCN	PMA.Lib	34816 x 60624	00:00:42	473.99 mb	10833	256.24	0.030
SVS	PMA.Lib	46000 x 32914	00:00:34	184.22 mb	7869	225.96	0.033
VSI	PMA.Lib	77231 x 130363	00:03:39	539.6 mb	51482	234.03	0.032
Huron	PMA.Lib	78490 x 94715	00:00:44	51322 mb	7510	168.1	0.049

