

SMEDGE

What's New in Smedge

Smedge 2012 Update 2

© 2004 - 2012 Überware™

Table of Contents

INTRODUCTION 4

UPDATE 2 5

- List Filtering now allows searching for any term, all terms, or an exact match
- When engines go to sleep, the status shows up as "Asleep"
- There is an option to keep a certain number of Engines awake at all times
- New environment variables to override the wake up process
- New Job parameter to override the maximum failure count for a specific Job
- New option to remote connect to Engines by IP address
- Fixed bugs in the power management system
- Engine and Job counts are now correctly reset when the GUI is disconnected
- Maxwell now limits the seed number to 16383
- Some components were missing from some distributions
- Master could get into infinite loop trying to kill zombie workers
- On Windows, fixed possible crash on startup if the first network interface from the OS did not have a MAC address
- On Linux and Mac, extra space characters in an executable path could cause a crash
- On Mac, some component processes were not being started internally with the correct path

UPDATE 1 9

- Immediate command execution on the Engine
- Engine disables itself if the disk is full
- Improved Power Saving Operations
- New GUI command to update a parameter for multiple Jobs
- Improved error message display
- Better hardware detection
- Job Finished Time is available for display in the GUI
- GUI displays a global count of objects in the status bar
- Submit command line shell "Update" command no longer needs to be quoted
- Aegis only starts a single instance per user
- New FFmpeg Virtual Module
- New REDline Virtual Module
- Maya products now show the command formatting string
- Maxwell deletes merged MXIs after merging to free up disk space
- Maxwell / Fry Modules now keep a running unique starting ID
- MAYA_LOCATION environment variable now set only if it is not already set
- The Mark Job As Finished command is now covered with the "Change Job" restriction

UPDATE 1 BUG FIXES 12

- Log cleanup can get denied
- Engine daemon not saving anything into the History.log file on Linux
- Process::Execute is expanding parameters with quotes incorrectly
- GUI Job lists are not sorting correctly
- Job creation time and Engine start time get modified
- Job history can get doubled when the Master switches
- Job History saved file could be corrupt after an unexpected Master switch
- Machine control scripts are missing from Windows distribution
- Possible GUI crash if job is deleted after selecting
- View Process Output from History control can be wrong
- The Info panel is not updating automatically
- Fixed some possible conditions that could hang the Master
- Maya module triggers access violation

OPERATION 14

- Wake-on-LAN
- Remote machine control
- Variable packet size
- Automatic font synchronization
- New system to translate the contents of scene files
- Limit jobs per creator
- New maximum job failures limit
- Processor time and peak memory usage are now reported in the Job History
- New error detection test: Minimum memory usage at start-up
- Average work time calculations by the Master now normalize for the packet size
- You can now disable reporting of detected images
- Ignored error messages are reported in the Job History
- New option to automatically reset job failures when a job is deleted
- Extensive performance optimizations throughout

INTERFACE **18**

- Customizable GUI window and improved filtering
- Controllable Display columns in ViewLists
- Improved GUI component process control for Master and Engine running normally and as services (and lock files)
- Individual dispatch reporting for specific Jobs and Engines
- Job shell list command can now filter the results and can report a count as well as listing the matched jobs
- Engine shell can set product options
- New Submit shell built-in Pool "My Computer"
- Submit shell can update individual parameters for a job
- Conspectus window can scroll if there are more engines than fit on the window
- You can now reset failure counts for offline engines

MODULES **21**

- New Product: V-Ray Standalone
- New Product: imgcvt
- Maya Module correctly detects new image output formats in Maya 2012
- Maya submit script improvements
- After Effects integration script improved
- Improved error detection for Cinema 4D renders
- Cinema 4D submit script
- 3D Studio Max submit script
- Nuke submit script
- Modo Module shows additional statistics and fixes

PLATFORM SPECIFIC CHANGES **23**

- RPM installer for Linux
- Linux and Mac have more appropriate file permissions
- Mac now has version info available
- Work around for bug in Mac runtime libraries that could cause hangs writing to the log files

API **24**

- New SequenceDistributor parameter \$(ActualPacketSize)
- New ProcessJob error controls
- New RenderJob Parameters
- Job class includes a scoped helper class to disable automatic parameter transformations
- New scoped Increment/Decrement template classes
- Updated to wxWidgets 2.8.12

MAJOR BUG FIXES **26**

- Fixed several possible crash and hang points
- Master was ignoring the -Database flag
- Master did not always honor requests to see the Dispatch.log file
- Mirror Masters ignored the request to shutdown correctly if the Mirror option was not set in the Master Options
- Mirror Masters did not always persistently mirror the original primary master
- Clients could refuse to connect to the Master
- Engine did not always update settings correctly when editing engine settings
- Editing a Job that is in progress could cause rendering of already completed frames
- GUI sometimes did not update the Info Panel when you selected an item
- GUI will not show percent done at 100% if the job is not yet finished
- Virtual Modules now correctly translate File and Dir type parameters

Introduction

Welcome to Smedge 2012!

The latest version of Smedge provides a whole new level of performance and reliability, and gives you more control over your rendering workflow than ever before.

Smedge 2012 does not communicate with prior versions of Smedge. We strongly recommend that you update your whole network at one time to avoid strange any problems with old versions left running.

Because the data structures have also changed to allow for the new features, like Wake-on-LAN and automatic scene file content translation, we also recommend that you back up any old data before upgrading. While the data can be upgraded automatically, if you ever want to go back to the old version of Smedge, the upgraded data will not be successfully read by the old versions.

As always, if you have any questions about new features or our future development plans, we encourage you to contact us.

Thanks, and see you in Smedge!



List Filtering now allows searching for any term, all terms, or an exact match

In the “Name” and “Creator” fields, you can now choose 3 different modes for how the filtering works. To select the mode, right click on the control, to bring up the context menu. “Exact Match” means that all characters, including space characters, are used to match the string. “Any Term” will divide the string using the space characters and look for any term in that string in the job name. For example, if you type “02 03” (with a space between the numbers), it will match “shot 02 version 01” “shot 03 version 05” and “shot 02 version 03”. In “All Terms” mode, the string will still be divided into terms using the space characters, but in order to match, all terms must exist in the matched name. In the example above “02 03” will only match the job “shot 02 version 03”.

The context menu is available for both the view filter at the top of the view tab and in the interface for making your own custom views. As with all other filtering criteria, the value you set when creating the view is separate from the value used in the filter's own display. The default mode is “All Terms”.

When engines go to sleep, the status shows up as “Asleep”

Now, the Engine list will show the status as “Asleep” instead of “Offline”, and will keep track of how long the machine has been in the low-power mode (in the “Uptime” column, which, yes, is a little weird but still seems like the best place for it). The engine's color will also be different than the normal offline color (note that colors can be changed in the List Item Colors dialog box).

There is an option to keep a certain number of Engines awake at all times

There is a new option on the Configure Master dialog box that allows you to specify that there are always a certain minimum number of Engines online at any given time. When the Master wakes up Engines for new work, it will always try to wake this many extra, so they can be immediately available for new work. By default, this value is 0 meaning that no extra Engines are kept online, but some farms may find that having at least a small number of machines always ready can be helpful.

New environment variables to override the wake up process

By default, Smedge sends UDP broadcast packets on port 9 with the magic packet value needed to power on a remote machine using Wake-on-LAN. If you need to override the port for some reason, you can set the SMEDGE_WAKE_PORT environment variable to an alternate port number to use. Even better, you can completely override the wake up process with your own custom command string. Set the SMEDGE_WAKE_COMMAND variable to the command line, using a simple \$(Name) variable substitution syntax to get the specifics for each engine. Besides engine data like the name and id, you can get the last known IP address and MAC address for the machine, as well as the network

broadcast for the network or subnet that it lives on. This allows you to integrate any third-party command line based tool for waking the machines, or to wrap the wake up process with other tasks, like logging in a database. See the environment variables chapter of the Administrator Manual for complete reference of how this new command line system works.

New Job parameter to override the maximum failure count for a specific Job

You can override the maximum number of failures that a Job can have before no more work is set from that Job as an advanced parameter of the Job. If you leave this value at the default (-1), then the job will be limited by the Master's maximum job failures limit. Otherwise, the value you specify will be used to determine if a job has had too many errors to send out work.

New option to remote connect to Engines by IP address

The VNC command string now interprets an additional possible parameter when setting up the command line to open the remote connection to the Engine. \$(Engine.IP) will be replaced with the IP address of the machine instead of the name. If the IP address is invalid or unknown, then it will fall back to using the name, and relying on DNS name resolution to find the remote machine.

Fixed bugs in the power management system

Fixed some issues in the way the power management (sleeping and waking) systems worked, introduced by the changes in Smedge 2012 Update 1. The power saving systems are now much more efficient and intelligent about how machines are woken, using all of the criteria of the job that is triggering the action, and being correctly aware of the job status, pool, failure counts, and all other distribution information. It should now work correctly on all platforms under all circumstances.

Engine and Job counts are now correctly reset when the GUI is disconnected

When the GUI disconnects, it was not resetting the count of the Jobs and Engines in the system, so when it reconnected, the counts would be off. It now correctly resets the counts so that when it reconnects, they are accurate.

Maxwell now limits the seed number to 16383

Maxwell limits the CPU ID seed value to 16383 (the maximum value for a 16 bit short, signed integer). Smedge now correctly respects this limit, and will overflow the CPU ID seed value it uses to 0 once this many work units have been distributed. Unless you are rendering more than this many repetitions of a single job, this change should not affect you, other than to avoid a failure from Maxwell.

Some components were missing from some distributions

Some virtual modules and support scripts were missing or named incorrectly in some distributions. These have been corrected so that every platform has access to every possible component and works consistently.

Master could get into infinite loop trying to kill zombie workers

When an engine disconnected with a pending work unit being sent to it, the master's "zombie" cleanup process that catches these lost work units could get into an infinite loop forever trying to kill and requeue the missing work. This affected Master performance and reliability. This is now fixed, so the Master performs better and handles these situations more cleanly.

On Windows, fixed possible crash on startup if the first network interface from the OS did not have a MAC address

The MAC address of the first network interface reported by the OS is what is used to generate the machine's unique ID, which defines it to the system and is also used to encrypt license codes. If the network interface that happened to be listed first by the OS did not include a MAC address for some reason, then this would cause a crash as the process started. Now it avoids this situation by searching for the first interface that has a MAC address. If no interfaces are found with valid MAC addresses, the process will still fail to start, but should report more carefully about what the actual problem is so that it is easier to resolve.

On Linux and Mac, extra space characters in an executable path could cause a crash

The changes made by the fix in Smedge 2012 Update 1 for handling quoted parameters better caused a new problem to emerge if the executable itself had a space in the path or filename. This would cause the Engine to crash, leading to work being "Lost" and requeued, and a loss of use of the Engine until the process was restarted. Though the problem could be worked around by setting up a path without spaces, now it again correctly handles paths with spaces internally.

On Mac, some component processes were not being started internally with the correct path

Because of the structure of OS-X application bundles, and the fact that some Smedge components are bundled and other are not, there were some processes that were not being started with the correct path to the executable. For example, the Aegis buttons were not correctly controlling the Engine status, because they were not starting the Engine command line shell with the correct path. This is now correct so all components work as expected, and can start other components as needed.

Immediate command execution on the Engine

There is a new system for running a command immediately on a machine. In the Engine Machine Control sub menu, there is a new command that allows users to enter a command string into a prompt window which is then sent to all of the selected engines. The Engines will immediately execute this command whether they have any currently running work or not. There is also a new Engine command line shell command to send the immediate command from a command line or script.

Engine disables itself if the disk is full

The engine can now detect when the system temp disk or log disk is getting full. When either disk (if they are not on the same disk) is getting low on space, the Engine will disable itself, and set the Engine's note field to show what happened, including the path that triggered the status change.

Improved Power Saving Operations

The Wake-on-LAN system has been improved to give a smarter response to the new jobs. It now only tries to wake Engines that would be allowed to work on any queued jobs. It is also now smarter about dealing with Engines that don't wake and about how often it tries to wake up a machine. The settings for how often the engines are woken is set per engine in the new Power Options tab of the Configure Engine window. As with every other Engine setting, you can adjust the settings for multiple engines at the same time by selecting them all before you open the window.

New GUI command to update a parameter for multiple Jobs

Select the menu command **Job > Set Parameter** which brings up a dialog for entering the name of the parameter you wish to change, and the value that you want to set it to. The parameter update will be applied to all of the jobs you selected when you opened this dialog. This makes it very quick and easy to adjust a setting for many jobs at the same time. Any Jobs you have selected that don't recognize the parameter by its name, will simply ignore the request.

Improved error message display

If a work unit finishes on error, it now displays the last detected error message as the note for the top level history for that work unit. This makes it easier to see the real reason for a failed work, instead of being confused by the final processing messages that may arrive after the error message (usually the "Peak memory usage" message from the new system that monitors the child process memory usage). More information is still available if you open the history element, as before.

Better hardware detection

Smedge now detects and provides socket, cores, and hyper-thread information in the engine details display. This information is detected by the hwloc library, an open source library for detecting processor information. On Windows and Linux, this library is linked as a shared library, and the shared library is included in the Smedge distribution. On Mac, the library is linked statically and no additional libraries or dependencies are required.

Job Finished Time is available for display in the GUI

The time when the last work unit from a job finished is now maintained as part of the job data, and is available in the Job lists in the GUI. When updating from older data, the last history element in the job provides the job finished time. Any jobs that finish when running the new version have the time marked by the master specifically in the job data. The column can be displayed by right clicking on the header of any job list in any view.

GUI displays a global count of objects in the status bar

The GUI will now show you a summary of statistics for the entire system, with the number of Engines, Jobs, and Work units currently active. It also displays the count of engines that are online and active, and the count of jobs that are finished and active. The statistics are global and shown in the status bar no matter which view is active.

Submit command line shell “Update” command no longer needs to be quoted

The Update command used to require that the value being passed was a single parameter. Values with spaces would have to be quoted. Now it just takes the remainder of the command line as the value. The following commands will now have identical results:

```
Submit Update Name "This is a name"  
Submit Update Name This is a name
```

Aegis only starts a single instance per user

Now, if it is already running, it won't open a second copy of the window.

New FFmpeg Virtual Module

A new Virtual Module that controls ffmpeg, allowing you to queue movie generation and extraction operations. Because of the sequential nature of most movie file operations, these jobs are submitted as a single work unit to one machine.

New REDline Virtual Module

A new Virtual Module that controls the REDline command line tool for converting RED camera footage.

Maya products now show the command formatting string

Maya products now allow you to override the command line formatting string in a similar way to Virtual Modules. The default command formatting string is still defined in the INI file, but it can be overridden in the Engine's Configure Product Options dialog. This allows you to customize how the Maya command line is used for a Maya based product without having to unload the module or shutdown the Engine, and copy the INI file around.

Maxwell deletes merged MXIs after merging to free up disk space

Since the merge process takes two input MXIs to generate a new output MXI, the input MXIs are no longer needed after the new file is generated. These old files are now deleted once they are merged. This will help reduce the disk load on the machine during processing of Maxwell renders. The MXI files are also all moved into the job local temp folder after downloading to ensure that they are all cleaned up properly.

Maxwell / Fry Modules now keep a running unique starting ID

Previously, each repetition work unit for these jobs were assigned an ID starting from 0 for each new job. This is now changed to use a unique serial number for its ID to help prevent conflicts. This means that you can resubmit a Maxwell or Fry job to get more repetitions and you will be able to merge the results together without having repeated IDs.

MAYA_LOCATION environment variable now set only if it is not already set

This environment variable is required on OS-X for Maya to function properly and is automatically set by Smedge based on the path to the executable. However, if you now set this variable yourself before starting Smedge, the value you set is respected and it no longer attempts to set the variable itself.

The Mark Job As Finished command is now covered with the “Change Job” restriction

Rather than the “Delete Job” restriction, since the job doesn't get deleted. This makes it more useful for people who may not have permission to delete any jobs but want to make it clear that a job is all finished so no more work is sent from it. Also removed the “Delete All Finished Jobs” command from the Job List context menu, because it is an operation that doesn't care about the selected job items.

Log cleanup can get denied

Fixed a problem in log cleanup which could cause logs to grow uncontrollably when under heavy load.

Engine daemon not saving anything into the History.log file on Linux

Fixed a bug in which the engine running as a daemon stops saving anything into the log file

Process::Execute is expanding parameters with quotes incorrectly

The process command line execution was not always running correctly as parameters in Unix environments were having their quotes and escaped quotes parsed incorrectly. This has since been fixed to use the correct parameter values when starting a child process.

GUI Job lists are not sorting correctly

The view columns were sorting on the wrong parameter values. Changed it to use the correct sorting methods. Also, the Job staggering now displays “No Staggering” if the system is disabled and sorts the staggering column based on both stagger settings more intelligently.

Job creation time and Engine start time get modified

When engines communicate times back and forth to the master, the local time adjustments could cause problems and alter the original time reported by jobs and engines. This has been resolved and the times are no longer drastically altered.

Job history can get doubled when the Master switches

When a mirror takes over as the primary Master, the job history could have duplicate entries for some of the recorded events. Fixed a problem with the messaging system and resolved the problem with duplicate history messages.

Job History saved file could be corrupt after an unexpected Master switch

The Job history was being reconstructed whenever a switch happens. However, the new master did not always have the correct job history in memory. Changed the design to be more reliable on unexpected master-mirror switches.

Machine control scripts are missing from Windows distribution

These files are now correctly included with the new release, in the Utilities folder.

Possible GUI crash if job is deleted after selecting

If the job is deleted after its selected but before the info pane is updated, a crash could occur. Improved the code to check for this situation and avoid the problem.

View Process Output from History control can be wrong

If work element has multiple attempts to run, selecting "View Process Output" when clicking on the un-expanded history for the work, should always show the output from the most recent attempt. However, it seemed to always show the output from the first attempt only. This problem has since been fixed to do the correct behavior.

The Info panel is not updating automatically

Some of the info panels, such as engine failures, were not being updated correctly. Fixed the idle event processing to ensure that the panels update correctly as needed.

Fixed some possible conditions that could hang the Master

There were a few conditions that could cause the Master to deadlock threads and appear to hang under heavy loads. These conditions have been corrected to avoid the deadlock conditions.

Maya module triggers access violation

Maya module triggered an access violation setting up command line for imgcvt Product, as imgcvt does not have some of the parameters that are assumed to exist by the environment variable setting method. The environment setting event handler now handles the lack of parameters correctly.

Wake-on-LAN

Smedge now includes a sophisticated system to allow you to reduce your farm's power consumption using the “Wake on LAN” technology built into most modern machines. After a period of idle time, Engines can now be set to automatically put themselves into their low power “sleep” mode. Then, when new jobs are submitted, the Master will wake up as many nodes as it needs to get the work complete. You can configure the timeout period for the machines to be put to sleep in the Engine Settings dialog.

The ability to wake machines is also available both from the GUI and using the Engine command line shell

Remote machine control

There are now several commands available in the Engine menu that allow you to shutdown or restart a machine, or put a machine into low power “sleep” mode. These commands all use a simple script system (batch files on Windows, and shell scripts on Linux/OS-X) that allow you to customize exactly how it works, and even configure special permissions that may not be available to the Engine itself. Additionally, there is a new command to force all Smedge component processes on a machine to shut down, making it possible to cleanly stop all Smedge on a remote machine (say if you want to perform an upgrade or configuration change) without having to go to the machine and find and kill those processes manually.

These commands are all available in the Engine menu, and are only available in “Administrator Mode.” The commands to sleep, restart or shutdown the machine do require that the Engine component is running and connected on the machine, however the commands to wake a machine up or to stop all Smedge component processes are available even if the engine is offline or disconnected (though they may have no effect if the configuration is incorrect or if the processes cannot communicate with the Master). These commands are also now available from the Engine command line shell.

Variable packet size

You can now set the Master to vary the packet size based on the Engine's “Priority” value. The system uses the priority of each engine as a percentage of the given packet size, allowing you to set up your machines to get more or less consistent render times even on widely varied hardware.

For example, if you set your Job packet size to 20 frames, and you have your fast machines set to priority 100 and a few slow machines set to priority 50, then the fast machines get the full 20 frame packet but the slow machines get 10 frame packets.

When this system is enabled, the Job History displays only the total and average packet times and does not attempt to determine “per frame” times (since each machine is rendering different numbers of frames).

Automatic font synchronization

Smedge can now automatically ensure that fonts from a specified folder are all correctly installed and available on a machine before starting work. This system is available for any product in Smedge, but is primarily useful for programs like After Effects that can make direct use of system fonts as part of their operation. The fonts are installed only if they are not available, and will be removed from the system after they are no longer needed, to ensure that your system does not get bogged down with an excessive number of fonts over time.

New system to translate the contents of scene files

There is a new Parameter Command called “TranslateContents”. If the value of the parameter is a valid path to a file, and if there are path translations set up on your Master, using this command will make a copy of the file in the Job local folder. Each line in the file will be searched for any translatable roots, and if found it will translate the path as it writes the file to the local folder. If the parameter doesn't correspond to a readable file, or if there are no translations configured, the command has no effect and the parameter value is untouched.

You can enable products to automatically use the TranslateContents command for the Scene parameter, using the new “TranslateSceneContents” parameter, set in Job's Advanced Info dialog under Render options. When enabled, any time you access the \$(Scene) parameter, it will be translated locally (if it has not already been) and

Limit jobs per creator

You can now set up a limit to how many workers are allowed from jobs with a specific Creator. This is another option for ensuring that your farm resources can be shared, without having to set up special Pools for every possible user. These limits are set in the Configure Master dialog box.

New maximum job failures limit

Smedge can now limit how many times a single job is allowed to fail before no more attempts are made to send out work from the Job. Before, you could limit how many times specific Engines were allowed to fail on a Job (and on a type of Job), but the Job itself may get passed around to every machine in your network. This could cause jobs to accrue hundreds of errors before they finally failed on every possible engine, and could cause delays in getting to other work caused by the failing job. Now, you can set a maximum number of failures for the Job itself regardless of the Engine failures, to keep these failing jobs from wasting too much time. This option is set with the other failure limits in the Configure Master dialog box.

Processor time and peak memory usage are now reported in the Job History

Smedge previously monitored the processor usage and total elapsed real time of a Job, but did not make that information readily available without using the API. Now, that information is directly reported in the Job's history. Additionally, Smedge now also monitors the process' peak memory usage and reports that as well. These systems all work consistently on all platforms.

Note that it is possible on the Unix based platforms (Linux and OS-X) for 3rd party applications to defeat this system. For example, After Effects essentially launches its own render daemon, and completely disconnects itself from the traditional Unix based process hierarchy. There is no way for Smedge to use only operating system resources to monitor these kinds of processes, so you may see incorrect statistics reported for AE renders on Macs until Adobe changes how their command line renderer works or provides some kind of public interface to find the spawned “daemon” process.

New error detection test: Minimum memory usage at start-up

Related to the memory monitoring system, Smedge can now ensure that a render process exceeds a certain amount of memory usage within a certain period of time after the process starts. This is another way to catch processes that can occasionally hit infinite loops during their initialization phase. For example, 3D Studio Max can occasionally hang when starting up, using CPU time but never actually starting the render. This system can now be used to detect these conditions, triggering an error and forcing the work to be restarted.

Average work time calculations by the Master now normalize for the packet size

When the Master is monitoring for packets that are out of the range of the average packet time, it will now normalize the time based on the packet size, so that 1 frame packet at the end does not get incorrectly detected as an error because it is so much slower than the other packets that are rendering 10 frames at a time.

You can now disable reporting of detected images

If you are not using the Smedge image detection system for anything in particular, you can now disable the reporting of the detected image files from the Engine back to the Master. The Engine will still do its normal error processing on the image files it detects, but the filenames will not appear in the job history. This can speed up the master and the GUI processing, especially if a render generates a lot of different images, but some Smedge features that use the image file names will no longer be available. You can set this option in the Configure Master dialog box.

Ignored error messages are reported in the Job History

When Smedge finds an error string that is set to be ignored, it now reports that message in the job history. This is useful so you can see if there are errors in the scene that may be being ignored but that should not. Also, if an error message is ignored, you can now optionally set the engine to

also ignore the process exit code. This is useful for renders that may produce benign errors that are correctly ignored, but that cause the rendering process to return an “unsuccessful” exit code anyway. If this option is enabled (in the Advanced Job Info) then the process exit code is only ignored if an error string is detected but ignored. If no error messages are detected, then the process exit code test works normally. There is also an option to disable reporting of the ignored errors, if you don't want to use the extra memory and bandwidth to report them in the history.

New option to automatically reset job failures when a job is deleted

In the Configure Master dialog you can now set the Master to automatically reset all failures from a job that gets deleted. Without this, once you delete a job the only way to reset the failures from that job is to reset the engine failure counts on every engine that failed on that job.

Extensive performance optimizations throughout

All Smedge components have been thoroughly profiled and optimized to make the system much more responsive and scalable.

Customizable GUI window and improved filtering

The views available in the GUI are now completely customizable. By using the menu command View > Customize Views, you get a new tab that allows you to create your own custom tabs in the interface. You can configure the default “filters” that are used for the views, and give them your own names, as well as change the order in which the views appear. There is also an option for removing views and for restoring the default views. For each view, you can select which list controls will be shown and use the filters to customize which items will populate the lists.

There are also more filtering options available now. Besides the name, creator and pool, you can now filter based on the job status and active products, and you can specify if you want to filter based on all of the options or any of them (AND or OR the filter items). The filtering you set up for your custom view here does not affect the filter that is also available at the top of each view itself, so you can still further filter each custom view to find specific items in each list.

See the User Manual for more information on the customization options and how it all works.

Controllable Display columns in ViewLists

You can now control the columns displayed in each view's Engine, Job, and Work ViewLists. You can right click the header / name of each column to bring up a context menu with some manipulation options. You can move the column you selected by right-clicking, left, right, to the extreme left, or to the extreme right. You can remove a column from being displayed by clicking on “Remove” or clicking on the name of that column if its checked as being displayed. If the column name in the menu is not checked, you can check it to bring that column display up. There's a restore defaults option to reset the column display to what smedge originally displays. Finally, there's a Custom columns option where you can define yourself, what's displayed with a parameter command string, in the form “\$(ParameterName).”

Improved GUI component process control for Master and Engine running normally and as services (and lock files)

The Components submenu of the System menu is now more clear and gives you more options for controlling the Master and Engine components, both as normal processes and as services. There is also a new command available to attempt to remove the “lock file” for these components. This is helpful when something has left a stale lock file that is blocking the component from starting correctly.

Individual dispatch reporting for specific Jobs and Engines

You can now get an individual dispatch report for specific Engines or Jobs to find out why that Engine is not starting work or that Job is not being dispatched. This can make it a lot easier to see what is going on, especially if you have a very large number of jobs or engines, and don't want to have to deal with the entire dispatch log, which can get quite large. To get the report, select an Engine or Job and choose the appropriate View Dispatch Report command from either the Job or Engine menu. These commands are also available in the context menus for the Job and Engine lists.

Job shell list command can now filter the results and can report a count as well as listing the matched jobs

The Job command line Shell executable has new options available when listing jobs. You can now supply similar filtering as the GUI list controls, to limit the results displayed to the jobs that fit your criteria. You can filter based on the name, creator, pool, status and type, with options to negate the filtering (i.e., all jobs except those named “foo”), and you can set up if the filtering requires for all items to be matched or any one of them.

Engine shell can set product options

There is a new Engine shell command that allows you to set individual product options by command line. For example, if you want to update the path to the executable for a specific product, you can supply the product name or ID, then the option name (e.g. “Executable”) and the value for the new option. Using the shell you can set the option for multiple machine simultaneously just like any other Engine operation.

New Submit shell built-in Pool “My Computer”

You can supply the pool name “My Computer” to the Submit shell, and it will set the pool for the Job it creates to the ID of the local machine. This is the machine that is running Submit, so this is most useful for scripts that integrate with a third party application, allowing end users to easily submit jobs that run only on their local machine, without them requiring to set up a pool manually or know the ID of their local machine.

Submit shell can update individual parameters for a job

You can now use Submit to change individual parameters for existing Jobs in the system. The new command is “Update”, after which you must supply the ID of the job to update, the name of the parameter to change, and the new value.

Conspectus window can scroll if there are more engines than fit on the window

Instead of becoming impossible to read, when there are more engines than can fit in the window, Conspectus will now make its display scroll so you can see every engine at a reasonable size.

You can now reset failure counts for offline engines

Because the failure counts are maintained by the Master, there is no reason not to allow resetting failures for Engines that are not actually online, so now you can.

New Product: V-Ray Standalone

There is a new Virtual Module to control the V-Ray standalone renderer.

New Product: imgcvt

There is a new Product defined by the Maya.sx module to control the imgcvt image conversion tool that is included in all Maya distributions. This module can spread the conversion of an image sequence across a farm making it very fast to convert from one format to another. It will also automatically detect the sequence specifiers and range of a sequence when you browse for files in the Smedge GUI.

Maya Module correctly detects new image output formats in Maya 2012

mental ray changed how it outputs the filenames in Maya 2012, and the Maya module has been modified in order to correctly detect image files from both Maya 2012 and earlier versions.

Maya submit script improvements

The smedgeRender.mel script that allows you to submit jobs to Smedge from directly inside of Maya has been greatly improved. It now provides nearly identical control as using the Smedge Gui to manually submit a job. It also now allows you to configure the path to the Submit executable yourself with a nice GUI interface in Maya instead of having to edit the script manually. You can download pools and select which pool to use, and it gives some advanced options, like first copying the scene file to the Maya renderScenes folder in your project to avoid possible issues as you keep working. Plus it now automatically integrates itself into the Render menu in the Maya GUI when you source the script.

The openPipeline version of the script is now also included in the Smedge distribution. This script is called op_SmedgeRender.mel and provides the same level of integration as the normal smedgeRender script, but is designed to work with the openPipeline project management system. To use this script, you should put it into the “addons” folder of your openPipeline installation, and it will be automatically available when you start openPipeline.

After Effects integration script improved

The After Effects submit script has also been significantly expanded. It is now designed to create a dockable panel that you can use as part of your AE interface to submit your project to Smedge. It allows you to download the pools from the master and select a pool, as well as configure several

other options for how the job will render. You can also now use a simple graphical UI from inside of AE to set up the path to the Submit executable, so you no longer need to manually edit the script file to get it to work correctly.

Improved error detection for Cinema 4D renders

Several new possible error messages were added to the Cinema 4D virtual module file to help aid in detecting when renders fail with Cinema 4D. Unfortunately the C4D command line renderer is quite terse, so even with the new messages it can be a bit of a challenge to figure out what is actually wrong, but at least more of the failures will be correctly detected as failures now.

Cinema 4D submit script

There is a new, very simple, submit script for submitting jobs to Smedge from inside of the Cinema 4D interface. This script is still quite simple, and you must manually specify the path to the Submit executable in the script for it to work correctly. It does not currently provide any interface.

3D Studio Max submit script

There is a new, very simple, submit script for submitting jobs to Smedge from inside of the 3D Studio Max interface. This script is still quite simple, and you must manually specify the path to the Submit executable in the script for it to work correctly. It does not currently provide any interface.

Nuke submit script

There is a new simple submit script for submitting jobs to Smedge from inside of the Nuke interface. This script is still quite simple, and you must manually specify the path to the Submit executable in the script for it to work correctly. The script provides a very simple interface for supplying some of the basic job parameters when you submit the job.

Modo Module shows additional statistics and fixes

Modo module now shows extended output information. It has been tested with Modo 601 and works fine. There is also a fix that allows you to correctly override the output path for your Modo renders.

Platform Specific Changes

RPM installer for Linux

Both the 32 bit and 64 bit builds of Smedge for Linux are now available as RPM files. The RPM installs Smedge into a standard location (`/opt/smedge`), and creates the machine folder (`/etc/smedge`) with appropriate permissions. The RPMs also handle the required dependencies for the Smedge component applications. Of course you can still download Smedge as a compressed TAR archive, and install it in a custom manner if you wish, but you will need to make the machine folder manually, and you will need to install any required packages yourself.

Linux and Mac have more appropriate file permissions

Though not technically a problem, the Unix based distributions should have more appropriate permissions on the component files. Only the executable files will have the executable bits set.

Mac now has version info available

The Mac release now displays the Smedge version information when you use the Get Info command in the Finder on the Smedge application icon.

Work around for bug in Mac runtime libraries that could cause hangs writing to the log files

Smedge now works around a bug in the Mac runtime libraries that could lead to a component hanging if it tried to write too much data to a file. The problem still exists (it is a bug in either the low level OS or the C runtime libraries provided by Apple), but Smedge should stumble into it only extremely rarely now.

New SequenceDistributor parameter **\$(ActualPacketSize)**

This new parameter in SequenceDistributor will report the actual number of frames in a packet, as opposed to the **\$(PacketSize)** parameter, which is the maximum number of frames in a single packet. Where **\$(PacketSize)** comes from the parent job and will be the same for all workers, this parameter will vary depending on how many frames are actually in a specific work unit.

New ProcessJob error controls

ProcessJob has added some more advanced error handling controls. When an error is ignored, by default it now reports the ignored error in the Job History. You can use the **\$(ReportIgnoredErrors)** parameter to view or adjust this behavior. Ignored errors will also now trigger Smedge not to examine the process exit code, as well, which is controlled by the parameter **\$(ExtendIgnoredErrors)**.

ProcessJob also now monitors the child process for the memory and processor usage, which it can report in the Job History. (See [here](#) for more information.) Though the values are not reported until the end of the work, Smedge can now monitor the child process to ensure that it exceeds a specifiable amount of memory usage within a certain period of time. This is useful for catching failed renders that may be caused by infinite loops in the rendering process at the start of rendering. (For example, 3DS Max can have this issue, and Smedge has some low level defaults that should catch this kind of failure most of the time). This can speed up detection of renders that will hang indefinitely without having to wait for the render to exceed the Overtime Kill ratio. The parameters that control this are **\$(StartupMemoryJob)** to set it as a Job parameter and **\$(StartupMemoryEngine)** to set it as an engine product option. The value is formatted as: *memory-in-mb/seconds*.

And you can access the result processor and memory usage events at the end of the work using **\$(ElapsedRealTime)**, **\$(ElapsedProcessTime)**, and **\$(PeakMemoryUsage)**.

New RenderJob Parameters

The new RenderJob functionality has associated new parameters to support it. To set or query if the image files are reported to the Master, use the **\$(ReportImages)** parameter. To set or query if the automatic scene contents translation system is enabled, use the **\$(TranslateSceneContents)** parameter. To set or query the [font synchronization folder](#), use **\$(FontSyncFolder)**.

Job class includes a scoped helper class to disable automatic parameter transformations

The new scene translation system is implemented using an automatic transformation applied to the parameter value when it is queried. In order to allow normal operation (where the raw value is returned), the API now provides a simple scoped class you can implement to disable the automatic

transformation from being applied. Any future automatic transformations that may be added will automatically be part of this system, future-proofing any code that makes use of these systems.

New scoped Increment/Decrement template classes

Several scoped classes have been handled to better handle incrementing or decrementing values in an exception safe manner.

Updated to wxWidgets 2.8.12

Smedge now uses wxWidgets version 2.8.12 (the current latest in the “stable” development branch). For best results, you should use this same version of wxWidgets if you want to create your own graphical components that make use of the API.

Fixed several possible crash and hang points

Crashes are a fact of life, but several causes of crashes and hangs in the Smedge code have been detected and resolved or avoided. This version of Smedge should be far more reliable, even under the heaviest of loads. With a system with this much open operation and this much dependence on the reliability of third-party applications, of course it is impossible to avoid all possible ways that you can crash the machine, but the code is greatly improved and Smedge itself should be the cause of these problems vary rarely.

Master was ignoring the -Database flag

The Master was ignoring the request for an alternate database location using the -Database flag.

Master did not always honor requests to see the Dispatch.log file

The Dispatch log was not always generated and there could be a delay between the request to view it and the actual generation of the file that would cause the GUI to report that no log was available. Now the log is always generated and available, though the contents are only actually filled in when a request comes to view them (to avoid the performance penalty of logging the dispatch loop). Note that the new Job and Engine dispatch reporting system uses the same logging messages as this log, but filters only the messages appropriate for specific Jobs and Engines, so it is probably a more useful tool than parsing the entire dispatch report, which can be quite verbose.

Mirror Masters ignored the request to shutdown correctly if the Mirror option was not set in the Master Options

This could cause component processes not to shut down correctly, leading to strange behavior if there was a problem with a master, and difficulty in updating Smedge on a machine (since program files were still in use). This should no longer happen.

Mirror Masters did not always persistently mirror the original primary master

Some of the Master data was not being correctly persisted by the mirror instances of the Master on other nodes. They would work correctly when taking over duties as the primary Master, but after a restart, some data could be lost, including Engine settings, license codes, and Pools. These issues have been corrected and mirrors should be complete accurate copies of all data both when taking over while running and after restarting.

Clients could refuse to connect to the Master

It was possible for the communication system to get into an invalid state where the client thought it was connected but the Master did not agree. The conditions that would lead to this have been corrected, so communication is more reliable and you should not get the random engines that drop out or the random GUI processes that refuse to connect to the master.

Engine did not always update settings correctly when editing engine settings

Some engine settings were not being correctly updated, due to an issue with the “default” engine system. This is now corrected, so all engine settings are correctly updated.

Editing a Job that is in progress could cause rendering of already completed frames

There was a bug in the way the job distributor data was being managed that would cause all status information to be discarded when a change was made to a job parameter. This would cause any completed progress to be ignored and could lead to repeated rendering of a job.

GUI sometimes did not update the Info Panel when you selected an item

The Info Panel could sometimes get out of sync with the selection in the list.

GUI will not show percent done at 100% if the job is not yet finished

This was a rounding error. If a Job was at a percentage complete of 95.5% or higher, this was being automatically rounded to 100%. This looked weird (100% but not complete?) so it is adjusted now to never exceed 99% if the job is not complete.

Virtual Modules now correctly translate File and Dir type parameters

Fixed an issue with the Virtual Module system that caused custom parameters of type Dir or File not to be translated correctly.