

Debug Overview

This is preliminary documentation and subject to change

Last revised: 15 August 2001

1. Introduction

This document describes the debugging services in the Common Language Runtime (CLR). The first part outlines the goals and requirements for the services. The second part provides an overview of the CLR debugging API. Lastly, the third part gives a picture of the architecture of the debugging services.

1.1 Design Goals

The design goals for the debugging services in the CLR are as follows:

- The debugging model should be on par with what the CLR provides. That is, the model of a program as viewed through the debugging services should be similar to that of a program as presented to the runtime for execution.
- The CLR debugging API is a low-level API. It should provide the minimal functionality necessary to 'get the job done'. The API is called by debugging packages, rather than by end-users directly. There is no need therefore to perform extensive argument checking, for example – instead, it is the caller's responsibility to perform such checks
- It should be usable by existing tools. A large investment has already been made in sophisticated tools and debugging environments and the CLR debugging services should fit in with these as easily as possible.
- Changes in the execution logic when debugging versus when not debugging should be minimized. The goal is to avoid the scenario where a user cannot reproduce some behavior when debugging is enabled, and is therefore forced to debug by more primitive means.

1.2 Further Information

[Debug Reference](#) – defines the details of the debugging API

[Enabling Profiling and Debugging](#) – describes how to control JIT-attach debugging

2. CLR Debugging Requirements

2.1 Scenarios

The CLR debugging services must operate in the following scenarios. Note that CLR does not necessarily have to enable each scenario, but it must at least inter-operate with current methods of supporting them.

Out-of-process debugging

Out-of process debugging is when the debugger is in a process other than the process being debugged. This reduces interactions between the debugger and the process being debugged, thus allowing a more accurate picture of the process.

The CLR debugging services will directly support out-of-process debugging. The services handle all communication for managed code debugging between the debugger and managed portions of the debuggee process.

Even though the CLR debugging API is used out-of-process, some of the debugging logic (e.g., thread synchronization) occurs in-process with the debuggee. This is an implementation detail that should be transparent to the debugger.

2.4.2.3 Exceptions

The CLR debugging services allow a debugger to catch both first chance and last (second) chance exceptions in managed code. The thrown object is available for inspection at each point.

Native exceptions in unmanaged code are not handled by the CLR (unless they propagate up to managed code.) However, the Win32 debugging services shared with the CLR debugging services can still be used to deal with unmanaged exceptions.

2.4.2.4 Program events

The CLR debugging services notify a debugger when many program events occur. These include process creation and exit, thread creation and exit, application domain creation and exit, assembly loading and unloading, module loading and unloading, and class loading and unloading. To insure good performance, class loading and unloading events can be disabled per module if desired. Events for class loading and unloading are disabled by default.

2.4.2.5 Thread control

The CLR debugging services provide APIs for suspending and resuming individual (managed) threads.

2.4.3 Examining Program state

The CLR debugging services provide a detailed means to inspect the parts of a process which are running managed code when the process is in a stopped state. A process can be inspected to get a list of physical threads.

A thread can be examined to find out its call stack. A thread's call stack is decomposed at two levels. First it is decomposed into *chains*. A chain is a contiguous logical call stack segment containing entirely managed or unmanaged stack frames. In addition, all managed call frames in a single chain share the same CLR context. If a cross-thread call has occurred within the thread, the chains provide information to allow a debugger to easily track logical call stacks across threads.

A chain can be either managed or unmanaged. Each managed chain can be further decomposed into single stack frames. Each stack frame represents one method invocation. Stack frames can be queried to obtain the code they are executing. A stack frame can also be queried to obtain arguments, the local variables, as well as the native registers.

An unmanaged chain is not composed of stack frames. Rather, an unmanaged chain simply provides the range of stack addresses that are due to unmanaged code. It is up to an unmanaged debugger to decode the unmanaged portion of the stack and provide a stack trace.

Note that CLR debugging services do not understand the concept of local variables as they exist in source code. It is up to the debugger to map between local variables and where they are allocated (either at a variable index for MSIL functions or in a register/stack location for managed native code.)

Access to global, class static, and thread local variables is also provided.

2.4.4 Modifying Program State

Although it is an inherently dangerous operation, the CLR debugging services allows a debugger to change the physical location of the instruction pointer during execution. The instruction pointer may be changed successfully under the following conditions:

- The target instruction pointer is at a sequence point. Sequence points represent statement boundaries.
- The current instruction pointer is at a sequence point.
- The target instruction pointer is not located in an exception filter.

- The target instruction pointer is not located in a catch block.
- The target instruction pointer is not located in a finally block.
- From within a catch block, the target instruction pointer is not located outside the catch block.
- The target instruction pointer must be within the same frame as the current instruction pointer.

Variables at the current instruction pointer location will be mapped to the variables at the target instruction pointer location. GC references at the target instruction pointer location will be properly initialized.

After the instruction pointer is changed, the services mark any cached stack information as invalid and refresh the information the next time it is needed. Debuggers that cache pointers to stack information (frames, chains, etc.) should refresh this information after changing the instruction pointer.

The debugger is also allowed to modify the data of a program when it is stopped. Local variables and arguments in a function that is currently running can be changed in a manner similar to how they are inspected. Also, fields of arrays and objects can be updated, as well as static fields and global variables.

2.4.5 Edit and Continue

Edit and Continue is a feature which makes it possible to be in the middle of a debugging session, edit source code, recompile the modified source, and continue the debugging session without having to rerun the executable from the start. From a functional perspective, Edit and Continue implies the ability to modify the code that is running in the debugger while preserving the rest of the run time state of the executable being debugged.

Note: Edit and Continue mode is not supported in first release of the product. However, the information in this section is retained, for information.

It is intended to support the following scenarios in a future release:

Compiler Support. This scenario allows a debugger process to request a snapshot of the metadata for the running process, and have this marshaled back to the compiler to use. The compiler will be able to use the metadata APIs to open the metadata in Edit and Continue mode. This means that no tokens will move. Hence, if the code generation is deterministic, no changes should be perceived except those made by the user. Adding new metadata and having it reflected in the running process is not part of this scenario.

Replace an Active Method. The programmer makes a change to a JIT-compiled, active method and the change is reflected. The debugger will support hijacking of the active method for the EE to know when it hits the top of the stack. The EE will perform the required stub work to swap out the method. The common allocator will throw away the old method and the JIT compiler will compile the new method.

Adding Static Data. The programmer adds a new text string for rdata inclusion and then updates the program to use ldstr. The same holds for other instructions. This requires the EE to return a valid RVA base for the snapshot call, recognize the new data, and put the new data into the scratch data area of the running process.

Adding Local Variables. New variables are added to an existing method. This requires a new standalone signature in the metadata describing the new variable. The EE and JIT compiler extend the stack frame for a method which gets a new variable (layout of the original stack must remain fixed – this is by design).

Add a Static Method to a Class, Call It. Add a static method to an existing class definition. This requires the metadata to be extended with the new method definition and the method layout to be changed to include knowledge of the new method.

Add non-Virtual Method to a Class, Call It. This is same as the previous scenario, but only allow for non-virtual, non-static method to be added to the class. This change requires the method table layout to have knowledge about the new method, and allow calls to the new method.

Add Virtual Method to a Class, Call It. This is same as the previous scenario, only allow for a virtual method to be added to the class. This change requires the method table layout to have knowledge about the

new method, and allow calls to the new method. Space must be reserved in the v-table layout to allow for the new method to be added.

Add Virtual Method to Parent Class, Call It. This is same as the previous scenario, only the new method is in a parent class for which an instantiated instance of a derived class exists. Extra space must be reserved in the parent class's v-table to allow up to some maximum number of new methods to be added this way. This change requires new metadata and changes to the EE for method table layout.

Add Pinvoke Method to a Class, Call It. This is same as the previous scenario,, only for adding an Pinvoke surrogate value with full metadata. This change requires new metadata and some support in the EE to fault in the new library (if not already there) and the new entry point.

Override Method to Hide. The programmer may add a method override in a derived class that hides the parent method. The EE doesn't track all existing call sites that are made out of date. In theory the source dependency rules would force those targets to be recompiled and hence compiled with the JIT compiler using the proper reference.

Remapping Breakpoints. Allow breakpoints to be set in a method, then force the method body to be changed while keeping the breakpoints. This requires the breakpoints to be removed from the old method and mapped into the new method. As for any case where the method body changes, the debugger also needs to track the new JIT sequence point map. The compiler provides a map from the old MSIL sequence points to the new MSIL sequence points.

Add a New Class. The programmer may add a new class definition to a module. This change requires new metadata and EE to add the class to the delay load list.

Add a New Field to a Class. The programmer may add a new field to an existing class definition, of which there may be active instances. The change involves extending the metadata to include the new field, having the EE include this extra data, and any required work to provide the JIT compiler access to the new field. There may also be GC work to track what could be a new reference field.

The CLR does not provide functionality to delete methods and fields from a class. If a debugger supports deletion of methods and variables, it is up to the debugger to ensure that the deleted methods and fields are not used. A debugger could "tombstone" a method by using Edit and Continue to replace the deleted method with a method that throws an exception when it is invoked.

Edit and Continue operations may only be performed when the CLR is in a sane state with all managed threads stopped at safe points.

JIT Restrictions

In general, under Edit and Continue, the JIT-compiler tries to produce code (and stack frame) which is independent of previous versions of the method, but is still upgradeable from previous frames.

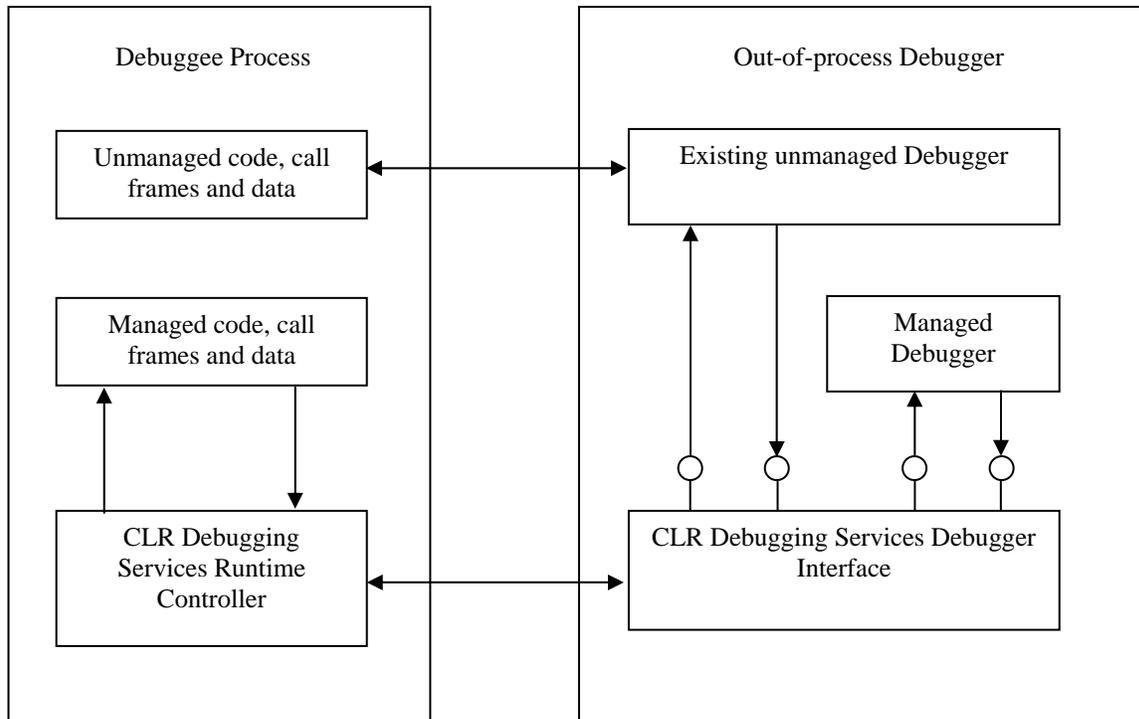
However, the frame layout depends on certain other factors and we don't put Edit and Continue restrictions on these. The JIT compiler has to handle these as gracefully as possible. All of these apply only to updating of a method currently on the stack.

- At the update point, the set of the older variables in scope has to be the same before and after. Newly added variables are acceptable.
- You can't be inside of a handler at the update point before or after.
- No Edit and Continue can be made if the older method used localloc.

2.4.6 Function evaluation

In order to evaluate user expressions and dynamic properties of objects, a debugger needs a way to run the code of the process being debugged. The CLR debugging services accomplish this by allowing the debugger to make a function or method call and have it run inside the debuggee's process.

Because this may be a dangerous operation (e.g., it may trigger a deadlock with existing code), the CLR allows such an operation to be aborted by the debugger. If the evaluation is aborted successfully, the thread will be treated as if the evaluation never happened, minus any side effects on locals, etc, from the partial



3.2 Managed Code Debugger

It is possible to build a purely managed code debugger. The CLR debugging services allow such a debugger to attach to a process on-demand using a “soft attach” mechanism. A debugger that is soft attached to a process can subsequently detach from the process.

3.3 Thread Synchronization

The CLR debugging services have conflicting requirements when it comes to process architecture. On one hand, there are many compelling reasons for keeping the debugging logic in the same process as the program being debugged. For example, the data structures are complex and in many cases are manipulated by APIs rather than by a fixed memory layout. It is much easier to call the APIs directly rather than trying to decode the data structures from outside the process. Another example to keep the debugging logic in the same process is for better performance. There is no need for cross process communication in this case. Lastly, an important feature of CLR debugging is the ability to run user code in process with the debuggee, which obviously requires some cooperation with the debuggee process.

On the other hand, CLR debugging must coexist with unmanaged native code debugging. This can only be performed correctly from an outside process. Also, an out-of-process debugger is safer than an in-process debugger because interference between the debugger’s operation and the process being debugged is minimized.

Because of these conflicting requirements, the CLR debugging services combine some of each approach. The primary debugging interface is out-of-process and coexists with the native Win32 debugging services. However, the CLR debugging services add the ability to *synchronize* with the debuggee process so that it can safely run code in the user process. To perform this synchronization the services collaborate with the OS and the CLR to suspend all threads in the process at a place where they are not in the middle of an operation which leaves the runtime in an incoherent state. The debugger is then able to run code in a special