

Chapter 1 : UDDI Version API Specification

The first part consisting of address bits a 0 and a 1 in the example in Figure) select the row. This selection remains active until revoked. This selection remains active until revoked. Then the second part, address bits a 2 and a 3, select the column.

SOAP fault reporting and fault codes will be returned for most invalid requests or any request where the intent of the caller cannot be determined. If any application level error occurs in processing a request message, a dispositionReport structure will be returned to the caller inside of a SOAP fault report. Faults that contain disposition reports contain error information that includes descriptions and typed keys that can be used to determine the cause of the error. Many of the API constructs defined in this document allow one or more of a given type of information to be passed. These API calls each conceptually represent a request on the part of the caller. The general error handling treatment recommended for UDDI operators is to detect errors in a request prior to processing the request. Any errors in the request detected will invalidate the entire request, and cause a dispositionReport to be generated within a SOAP Fault structure see appendix A. The error codes specified within each API call description are characteristic of the API call, but other UDDI error codes may be returned in unusual circumstances or when doing so adds additional descriptive information. White space characters include carriage returns, line feeds, spaces, and tabs. For this reason, the default collation order for data registered within an Operator Site is binary even though this choice is meaningless for some languages, and effectively favors alphabetic languages. Similarly, XML allows for a large number of character set encoding choices. The special values within API syntax examples are shown in Italics. In most cases, the following reference applies to these values: The name of the element or attribute designates the particular key type that is required. These keys are always formatted according to an algorithm that is agreed upon by the UDDI Operator Council with the one exception being tModelKey values, which are prefixed with a URN qualifier in the format "uuid: This special attribute is a required metadata element for all messages. It is used to designate the specification version used to format the SOAP message. As of the date this specification, any other value e. This special element is found in the inquiry API functions that are used to search i. This passed argument is used to signal special behaviors to be used with searching. See the findQualifiers appendix and the documentation for the individual find API messages for more information. This special qualifier is found in the inquiry API functions that are used to search e. This argument is used to limit the number of results returned from a request. When an Operator Site or compatible instance returns data in response to a request that contains this caller-supplied limiting argument, the number of main result elements will not exceed the integer value passed. If a result set is truncated as a result of applying this limit, or if a result set is truncated because the search would otherwise exceed an operator-specific limit, the result will include the truncated attribute with a value of true. The truncated attribute indicates that the results returned do not represent the entire query result set. The actual limit set for applying this treatment is Operator Site policy specific, but in general should be a sufficiently large number so as to not normally be an issue. No behaviors such as paging mechanisms are defined for retrieving more data after a truncated limit. The intent is to support the average query, while at the same time allowing Operator Sites the leeway required to be able to manage adequate performance. UDDI is not designed to support large data sets required by some research uses. Searches can be performed based on a cross section of categories. Several categories are broadly supported by all Operator Sites and provide three categorization dimensions. These are industry type, product or service type, and geography. Searches involving category information can be combined to cross multiple dimensions [9]. For this reason, these searches are performed by default matching on ALL of the categories supplied e. In general, the embedded category information serves as voluntary hints that depend on how the registering party has categorized themselves, but not to provide a full third party categorization facility. Searches involving identifiers are performed matching on any supplied identifier e. These searches allow broad identity matching by returning a match when any keyedReference set used to search identifiers matches a registered identifier. Version 2 provides for the definition of checked identifiers. This enhancement makes it possible to distinguish

copycat information within UDDI from the registrations of the authentic business registration based on validated identifiers. Searches that match a particular technical fingerprint use UUID values to search for bindingTemplates with matching tModelKey value sets. When used to search for web services e. For instance, the existence of a web service that implements all of the parts of the UDDI specifications can be accomplished by searching for a combination of tModel key values that correspond to the full set of specifications the UDDI specification, for instance, is divided into at least 3 different, separately deployable tModels. At the same time, limiting the number of tModelKey values passed in a search can perform broader searches that look for any web service that implements a specific sub-part of the full specification. For a full understanding of structure contents, refer to this schema as well as the UDDI data structure reference. It is suggested that tools that understand schemas be used to generate logic that populates the structures used to make the API calls against UDDI.

Chapter 2 : Best quality Kess Kess V2 V V ECU Programmer Online Version

Related products for Scania Diagnos & Programmer 3 (SDP3) v no Dongle + xcom Full Truck & Bus Package Online via RDP Online parts catalogues office for trucks is the best solution for workshops or spare parts sellers.

What every programmer should know about memory, Part 1 LWN. Without subscribers, LWN would simply not exist. Ulrich Drepper recently approached us asking if we would be interested in publishing a lengthy document he had written on how memory and software interact. We did not have to look at the text for long to realize that it would be of interest to many LWN readers. Memory usage is often the determining factor in how software performs, but good information on how to avoid memory bottlenecks is hard to find. This series of articles should change that situation. The original document prints out at over pages. We will be splitting it into about seven segments, each run weeks after its predecessor. Once the entire series is out, Ulrich will be releasing the full text. Reformatting the text from the original LaTeX has been a bit of a challenge, but the results, hopefully, will be good. Hyperlinked cross-references and [bibliography references] will not be possible until the full series is published. Many thanks to Ulrich for allowing LWN to publish this material; we hope that it will lead to more memory-efficient software across our systems in the near future. The various components of a system, such as the CPU, memory, mass storage, and network interfaces, were developed together and, as a result, were quite balanced in their performance. For example, the memory and network interfaces were not much faster than the CPU at providing data. This situation changed once the basic structure of computers stabilized and hardware developers concentrated on optimizing individual subsystems. Suddenly the performance of some components of the computer fell significantly behind and bottlenecks developed. This was especially true for mass storage and memory subsystems which, for cost reasons, improved more slowly relative to other components. The slowness of mass storage has mostly been dealt with using software techniques: Cache storage was added to the storage devices themselves, which requires no changes in the operating system to increase performance. Unlike storage subsystems, removing the main memory as a bottleneck has proven much more difficult and almost all solutions require changes to the hardware. Today these changes mainly come in the following forms: RAM hardware design speed and parallelism. Direct memory access DMA for devices. For the most part, this document will deal with CPU caches and some effects of memory controller design. In the process of exploring these topics, we will explore DMA and bring it into the larger picture. This is a prerequisite to understanding the problems and the limitations of efficiently using memory subsystems. We will also learn about, in some detail, the different types of RAM and illustrate why these differences still exist. This document is in no way all inclusive and final. It is limited to commodity hardware and further limited to a subset of that hardware. Also, many topics will be discussed in just enough detail for the goals of this paper. For such topics, readers are recommended to find more detailed documentation. When it comes to operating-system-specific details and solutions, the text exclusively describes Linux. At no time will it contain any information about other OSes. The author has no interest in discussing the implications for other OSes. One last comment before the start. The technology discussed here exists in many, many variations in the real world and this paper only addresses the most common, mainstream versions. It is rare that absolute statements can be made about this technology, thus the qualifiers. It does not go into enough technical details of the hardware to be useful for hardware-oriented readers. But before we can go into the practical information for developers a lot of groundwork must be laid. To that end, the second section describes random-access memory RAM in technical detail. Appropriate back references to the section are added in places where the content is required so that the anxious reader could skip most of this section at first. The third section goes into a lot of details of CPU cache behavior. Graphs have been used to keep the text from being as dry as it would otherwise be. This content is essential for an understanding of the rest of the document. This is also required groundwork for the rest. Section 6 is the central section of this paper. The very impatient reader could start with this section and, if necessary, go back to the earlier sections to freshen up the knowledge of the underlying technology. Section 7 introduces tools which can help the programmer do a better job. Even with a complete understanding of the technology it is far

from obvious where in a non-trivial software project the problems are. Some tools are necessary. In section 8 we finally give an outlook of technology which can be expected in the near future or which might just simply be good to have. This includes updates made necessary by advances in technology but also to correct mistakes. Readers willing to report problems are encouraged to send email. Markus Armbruster provided a lot of valuable input on problems and omissions in the text. Scaling these days is most often achieved horizontally instead of vertically, meaning today it is more cost-effective to use many smaller, connected commodity computers instead of a few really large and exceptionally fast and expensive systems. This is the case because fast and inexpensive network hardware is widely available. There are still situations where the large specialized systems have their place and these systems still provide a business opportunity, but the overall market is dwarfed by the commodity hardware market. Bigger machines will be supported, but the quad socket, quad CPU core case is currently thought to be the sweet spot and most optimizations are targeted for such machines. Large differences exist in the structure of commodity computers. Note that these technical details tend to change rapidly, so the reader is advised to take the date of this writing into account. Over the years the personal computers and smaller servers standardized on a chipset with two parts: The Northbridge contains, among other things, the memory controller, and its implementation determines the type of RAM chips used for the computer. To reach all other system devices, the Northbridge must communicate with the Southbridge. Older systems had AGP slots which were attached to the Northbridge. This was done for performance reasons related to insufficiently fast connections between the Northbridge and Southbridge. Such a system structure has a number of noteworthy consequences: All data communication from one CPU to another must travel over the same bus used to communicate with the Northbridge. All communication with RAM must pass through the Northbridge. The RAM has only a single port. It can be found in specialized hardware such as network routers which depend on utmost speed. A couple of bottlenecks are immediately apparent in this design. One such bottleneck involves access to RAM for devices. In the earliest days of the PC, all communication with devices on either bridge had to pass through the CPU, negatively impacting overall system performance. To work around this problem some devices became capable of direct memory access DMA. Today all high-performance devices attached to any of the buses can utilize DMA. This problem, therefore, must to be taken into account. A second bottleneck involves the bus from the Northbridge to the RAM. The exact details of the bus depend on the memory types deployed. On older systems there is only one bus to all the RAM chips, so parallel access is not possible. The Northbridge interleaves memory access across the channels. With limited bandwidth available, it is important to schedule memory access in ways that minimize delays. As we will see, processors are much faster and must wait to access memory, despite the use of CPU caches. If multiple hyper-threads, cores, or processors access memory at the same time, the wait times for memory access are even longer. This is also true for DMA operations. There is more to accessing memory than concurrency, however. Access patterns themselves also greatly influence the performance of the memory subsystem, especially with multiple memory channels. Refer to Section 2. On some more expensive systems, the Northbridge does not actually contain the memory controller. Instead the Northbridge can be connected to a number of external memory controllers in the following example, four of them. Northbridge with External Controllers The advantage of this architecture is that more than one memory bus exists and therefore total bandwidth increases. This design also supports more memory. Concurrent memory access patterns reduce delays by simultaneously accessing different memory banks. This is especially true when multiple processors are directly connected to the Northbridge, as in Figure 2. For such a design, the primary limitation is the internal bandwidth of the Northbridge, which is phenomenal for this architecture from Intel. Integrated Memory Controller With an architecture like this there are as many memory banks available as there are processors. On a quad-CPU machine the memory bandwidth is quadrupled without the need for a complicated Northbridge with enormous bandwidth. Having a memory controller integrated into the CPU has some additional advantages; we will not dig deeper into this technology here. There are disadvantages to this architecture, too. First of all, because the machine still has to make all the memory of the system accessible to all processors, the memory is not uniform anymore hence the name NUMA - Non-Uniform Memory Architecture - for such an architecture.

DOWNLOAD PDF V. 2 PROGRAMMER INFORMATION (PART 1 AND 2).

Chapter 3 : ST-LINK/V2 - ST-LINK/V2 in-circuit debugger/programmer for STM8 and STM32 - STMicroelectronics

Document Overview. This document describes the programming interface and expected behaviors of all instances of the Universal Description, Discovery and Integration (UDDI) registry.

The name of the module in which the error occurred. Member name if specified. The values of hex return codes and their meanings are as follows: Return Code Explanation 04 A data set opened for output used all space available to or on the current volume, and no more volumes were available. Change the JCL to specify more volumes. When all the volumes were filled, the program attempted to write another record. For a partitioned data set on a direct access volume or for a VIO data set, all space was filled when the program attempted to write another record. A partitioned data set or a VIO data set can reside on only one volume with a maximum of tracks. For a partitioned data set on a direct access volume, 16 extents had been used when the program attempted to write another record. All space was filled on the volume, and an attempt was made to obtain space on the next specified volume. Either the space was not available on that volume, the data set already existed on that volume, or there is no space available in the VTOC or the VTOC index. The message contains the volume serial number of the last volume used. If the error is to be ignored, the system will attempt to close the DCB before returning to the user. Operator response Start a generalized trace facility GTF trace, and re-create the problem. Reply to message AHLA with: Search problem reporting data bases for a fix for the problem. Programmer response Probable user error. Correct the errors causing the abnormal end as indicated by the return code in the message text as follows: For return code 04, case 1, specify at least one more volume than the number of volumes previously used for the data set. For case 2, specify a different volume for the partitioned data set or specify more space for the VIO data set. For case 3, either specify a volume for the data set, use a utility program to reorganize the volume so that data sets will not be fragmented that is, no more than 16 extents used for this data set , or change the program so that a device will be free when a volume must be mounted. For return code 08, either specify a new volume to continue the data set or make sure that enough space is available on the volumes already specified. Ensure that the data set does not already exist on the volumes to be used. In all cases, rerun the job. For return code 0C, consult your installation procedures.

Chapter 4 : CADILLAC CTS V L GAS QUICKTUNE PERFORMANCE & ECONOMY PROGRAMMER CHIEF

DRAGZINE: Rebuilding The Quadrajet Carburetor. Yes, They Are Worth It. A full and in-depth article of the rebuild process that Jet Performance takes with a Q-Jet Carburetor, by Dragzine.

Chapter 5 : Programmer - TV 2

No thanks 1 month free. Part one of information for part 2 luna Avery. Loading Unsubscribe from luna Avery? Cancel Unsubscribe. Working Subscribe Subscribed Unsubscribe 6.

Chapter 6 : Pololu USB AVR Programmer v

DiabloSport Predator 2 programmer is designed to increase your power and fuel economy, while maintaining a copy of your trucks factory data for future restoration. The Predator taps into your rides internal computer through the OBD-II port and adjusts it to increase horsepower and torque.

Chapter 7 : What every programmer should know about memory, Part 1 [www.nxgvision.com]

Â· Connection to PC: USB (up to Mbit/s, high speed/full speed) and compatible interface. Â· Comfortable and easy to use control program, works with all versions of MS Windows.

Chapter 8 : Scania SDP3 (Diagnos & Programmer) v with activation Download

The v programmer has improved circuitry for measuring VCC which limits the duty cycle of this effect to about %, so the motor won't move (but it might make a 25 Hz clicking sound). The v2 programmer would typically brown-out if a 5 V signal was applied to its RST pin while it was operating at V.

Chapter 9 : PRO-2 Handheld Programmer - Pyramid Technologies, Inc.

For return code 04, case 1, specify at least one more volume than the number of volumes previously used for the data set. For case 2, specify a different volume for the partitioned data set or specify more space for the VIO data set.