



RDM_e dataFlow with Raima Databases

A Raima Business White Paper

Published: February, 2010
Author: Duncan Bates
VP of Product Management

Copyright: Raima, a Birdstep Technology Company

Abstract

Today, the need for real-time data throughout the many levels of an organization is growing rapidly. The more current a company's business intelligence information often determines how successful a business is run. Traditionally, embedded applications have been closed systems. The movement of any data collected on these devices to the company's enterprise business systems is performed via off line batch processing at periodic intervals. This paper outlines three separate business cases applicable within three key industry market segments: Industrial Automation; Telecommunication; and Aerospace and Defense; where Raima technologies facilitate high speed data management and seamless real-time data flow between embedded devices and corporate enterprise databases. Additionally, this paper will outline (currently available upon request) a demonstration that highlights the use of the RDM embedded databases and dataFlow technologies on real-time embedded computers (devices), desktop and servers platforms. The demonstration relies on Raima technology partners' technology from Microsoft[®], Wind River[®], and Green Hills[®] Software.

This article is relative to the following versions of RDM:

- ✓ RDM Embedded: 9.1
- ✓ RDM Server: 8.2, 8.3

Table of Contents

Introduction.....	3
RDM _e dataFlow™ Software Components.....	4
RDM Embedded™.....	4
microHTTP™.....	4
Replication Server and Replication Client.....	5
RDM Server and Transaction Plug-Ins.....	5
Industrial Automation Use Case.....	5
Telecommunication Use Case.....	6
Aerospace & Defense (A&D) Use Case.....	7
The Live Building Power Monitoring Demo.....	8
Conclusion.....	10
Contact Information.....	10

Introduction

Today's complex computing environments are generating new requirements for data management and data availability that traditional solutions simply can no longer fulfill. Traditional embedded system solutions relied on very specific proprietary home grown data management code on the controller (device). If data was in fact ever moved from a device, it was most likely done so by uploading the data to a server in batches at some arbitrarily predetermined schedule. The servers in turn would distribute the data upward to corporate enterprises making the data available for business intelligence systems to make use of.

Systems today need to provide data synchronization, aggregation and distribution at all levels of an organization in order for an organization to be successful and meet its corporate obligations and responsibilities. Each level and discipline of an organization has diverse and often very different responsibilities that come together to meet the overall corporate goal of an organization. For example, at the shop floor level, the processing units require continuous operation. The Operations group has the responsibility of maintaining the system's uptime and corporate managers are tasked with monitoring and ensuring business performance. These are just a few of the many responsibilities throughout an organization.

Critical to any business's success and financial savings is the ability to catch system failures or unexpected behavior as early as possible. This is closely followed by the requirement to have the agility within the system to easily perform reconfiguration and maintenance operations. To best achieve these goals real-time access to data is key to the solution. The sooner a technician can identify a failing device the quicker he can take action to make repairs and keep downtime to a minimum or prevent it all together. The availability of real-time trend analysis and business intelligence enables management to implement corrective action to maintain its success.

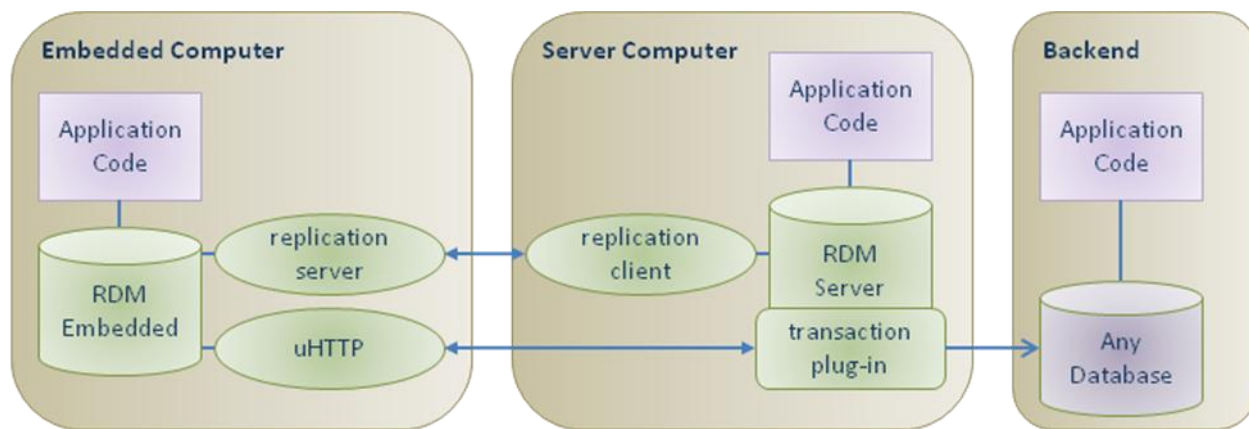
With all the diversity of runtime environments, requirements to data availability and scalability, and operational requirements at the different levels of the organization, a single data management solution simply cannot fulfill the requirements. Data management needs to be done at all levels with a diverse technology suited for the environment and requirements it's addressing. But safe data management is not the only answer; one needs an efficient, safe, and seamless way of moving data between these systems to solve application requirements at the different organizational levels.

With RDM data management solutions and the recently announced dataflow technology, you can safely manage data from the most resource constrained device up to the more robust enterprise servers. In addition, RDM technology will seamlessly replicate and aggregate your data between the various environments. All this is accomplished in near real-time, integrating the managed data into your corporate backend making business intelligence data available throughout the organization.

The rest of this document will outline the software components provided by Raima to achieve the aforementioned goals. This paper goes on to outline three real life use cases and concludes with the introduction of RDM dataFlow demonstration which is available upon request.

RDM_e dataFlow™ Software Components

The RDM embedded database technology has been in the embedded systems market since 1984 achieving great success and recognition in the managing of data on everything from the smallest device to the largest server. Today, you will find the RDM line of embedded databases in business critical systems of major industrial automation companies, some of the largest telecommunication equipment providers, and mission critical aerospace and defense contractors' applications. Now with the latest advances in our RDM products our developers of embedded systems can take the next step in solving complex distributed data management tasks. The diagram below illustrates how the Raima components are implemented to solve these problems.



RDM Embedded™

RDM Embedded is the core data management component which was developed specifically for harsh and resource constrained environments. It is available on a wide variety of software and hardware platforms ranging from the smallest devices that run without requiring any operating system up to enterprise servers running full blown server operating systems. RDM Embedded's role in this solution is to provide consistent, concurrent and safe access to well defined application data structures. The replication engine, of course, will recover from any system or application crashes and can be run on diskless as well as on disk enabled systems.

microHTTP™

Raima's RDM Embedded microHTTP server enables industrial standard access to the application database residing on the embedded computers. Adhering to the W3C standards this lightweight HTTP server can be run in both threaded and non-threaded environments to provide remote access to the application's data. The microHTTP server is tightly coupled with the embedded database but can easily be extended to web enable other parts of the application. The role of the microHTTP server is to enable real-time configuration changes and remote access to the embedded devices in the system solution. When coupled with RDM Server's transaction plug-ins, changes to server databases can be programmed to trigger real-time changes to the systems embedded devices.

Replication Server and Replication Client

Raima's RDM replication engine is comprised of two primary components; a server and client(s). The purpose of the replication engine is to safely replicate and aggregate data from one or more embedded devices to an instance of RDM Server running on a separate hardware and software platform within the system. Data is replicated in near real-time and can withstand any network interruptions because of its ability to re-submit any pending transactions. The replication server also makes optimal use of the network's bandwidth through configurable water level triggering to make the decision between full database copies or incremental transaction log replication.

RDM Server and Transaction Plug-Ins

RDM Server is Raima's client/server offering which within this solution acts as the aggregation point for multiple RDM Embedded instances (devices). The aggregation is performed in real-time and the results are made available for consumption to the application through standard SQL92. This provides access through high level API's like ODBC 3.51, ADO.Net, JDBC and other standard based interfaces. RDM Servers' sophisticated design includes the ability to plug application code (server extension) directly within its runtime. In this solution, this capability is utilized to provide real-time data feeds to the enterprise databases and the connected embedded devices. RDM Server is available on all the major desktop and server operating systems as well as a majority of the real-time operating systems on the market today. This gives developers the freedom to choose the environment best suited for their solution.

Industrial Automation Use Case

Raima's RDM technology is implemented in numerous industrial automations systems, ranging from complex turnkey power plant systems to simple batch processing systems. Common across many of these systems



is the requirement for real-time data management performance at the controller (device) level with automated data movement to upstream shop floor management systems and beyond to corporate management systems. One such application can be found within a petroleum processing plant. It's not uncommon to find systems today where the reliance on human intervention remains strong. Many of these systems are still manually monitored and sensor generated alarms are passed on to responders in the form of a phone call or radio communication.

For the most part, many of these systems were designed at a time when data management and dataflow solutions were still in their infancy. Despite this, one can just imagine the risks associated with these antiquated systems. Alarms are frequently delayed in their delivery to responders and many of them are in-accurate or altogether false. This may result in frequent and unnecessary maintenance, costly production stoppage or, in rare cases, catastrophic failure. Information about minor incidents seldom finds its way to top management, resulting in inaccurate business intelligence leading to bad business decisions. In this environment, the ability to make pro-active decisions is nearly impossible. Rather, because of the lack of real-time aggregated data system, operators wait for scheduled batch updates before re-configuration decisions can be made postponing tasks that could have increased the overall production.

The benefit of a robust data management and flow solution becomes obvious. With accuracy in events and alarms, the operational system will be much less prone to unnecessary maintenance and production stoppage. With real-time decision making, system operators can optimize total production and reduce risk by reducing the reliance on human capital. Business managers will, at any given time, have up-to-date information regarding the state of the system and accurate and reliable data for reporting and trend analysis.



Telecommunication Use Case

Ever since the introduction and following successive adoption of Raima's embedded database technology in the embedded systems market, the major telecom equipment providers like Alcatel-Lucent, 3Com and many others have utilized Raima's RDM real-time embedded database in their equipment. Adoption of RDM embedded database technology started with standalone voice and data switches and soon grew as advancements were made to these products. As the requirements for these systems expanded so did the capabilities of Raima's RDM embedded databases. Requirements that once were "nice to have" features have developed into standard requirements. Examples of these requirements are 24x7 operations with system failover capabilities, distributed data across multiple processing units, and real-time and in-memory data processing needed to keep up with increasingly faster network speeds.

One such application use case is a distributed IP intrusion detection system. The sole purpose of such a system is to capture IP packages and validate them against known patterns, and to shut down and prevent unwanted network traffic. One of the challenges of today's solutions is to implement a solution that can keep up with the data speeds commonly found in modern networks. Not only does the application need to log at the speed of the network, it also needs to do complex intrusion validation on a wide range of IP packages. Additionally, it needs to trigger real-time alarms to network administrators. Finally, it needs to notify other IP units about black listed traffic origins.

The main reason is that most of today's applications implementing proprietary home grown data management solutions are a direct result of the logging speed requirement. System developers having been forced into the decision making has resulted in the proliferation of proprietary alarming mechanisms and data distribution solutions. These systems are often far from ideal due to the absence of true transaction support within their data management implementations. Because of this shortcoming, it's common to find intrusion alarms being delayed, alarms typically stop at the network administrators, and distributed black lists are being updated based on a scheduled batch process instead of updating in real-time. Adding to the complication is the lack of transaction support results in the need to merge blacklists between units followed by a complete destructive refresh list. This limitation exposes serious vulnerability of the system which could potentially let through network traffic of whose origin had already been detected as hazards at one unit, but not at another.

With Raima's RDM embedded database true in-memory and circular data buffering IP intrusion units can efficiently add dynamic transactional safe data logging with real-time pattern matching. This combined with the dataflow capabilities described previously in this document, real-time alarms can be pushed upstream to the network administrators with ease, and distributed black list updates can be performed at a transactional level. One possible extension to this is the capability of alarms being pushed further up the system to management for on-the-fly reporting and trends analysis.



Aerospace & Defense (A&D) Use Case

Another strong industry vertical market for Raima is the A&D. Raima's embedded databases are found in a wide range of mission critical systems including ground to air communication systems, flight planning systems, radio tower systems, airborne data gathering systems, and hostile vehicle identification systems.

As with many applications found in the various industries the A&D market have very familiar requirements. Requirements such as the requirement for up-to-date information, real-time data exchange, system reliability, and security are common in A&D as they are in other industries.

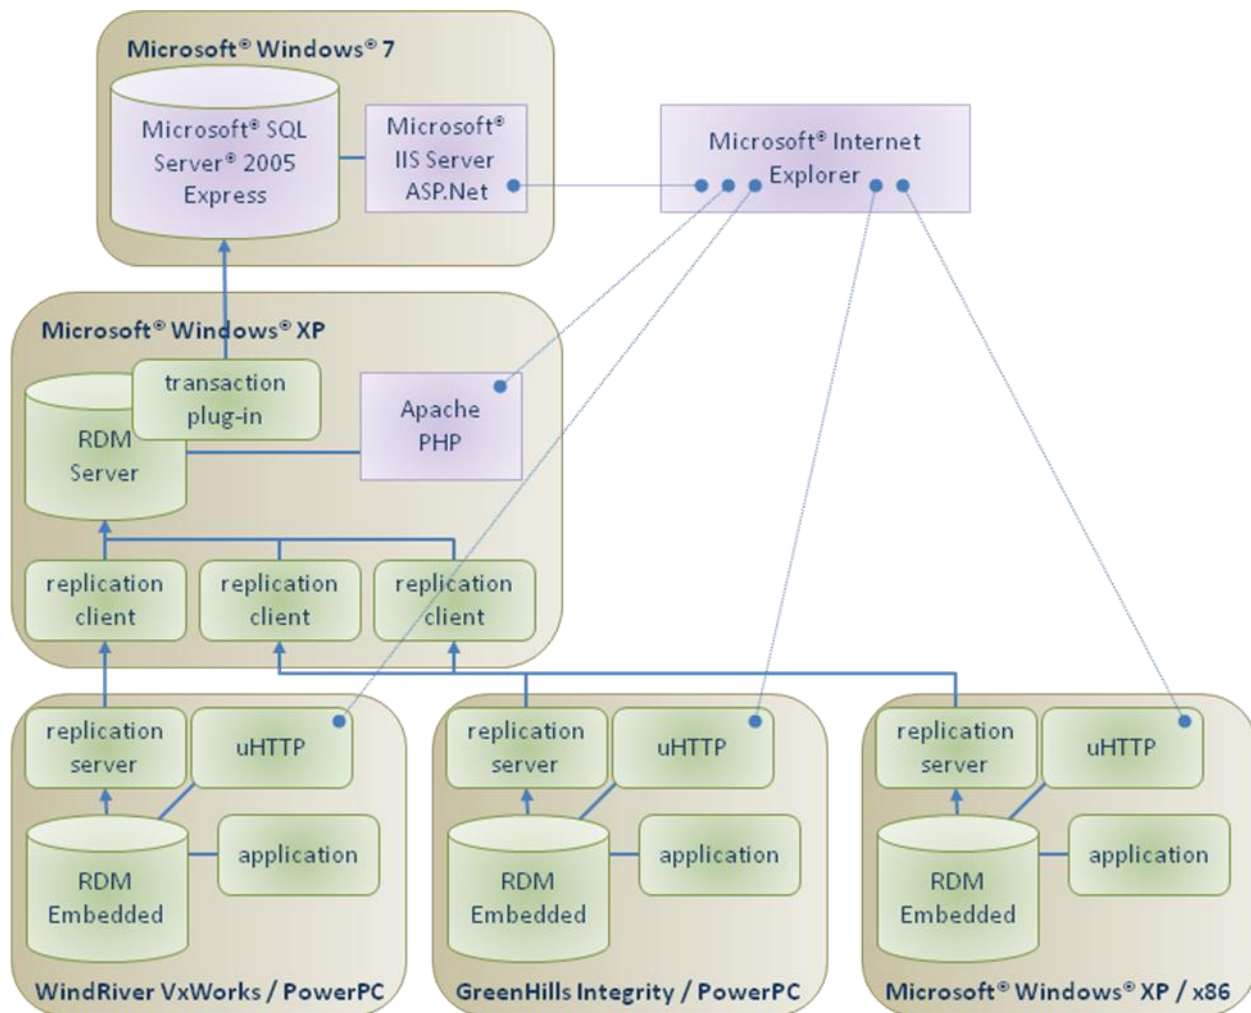
It's not uncommon to find systems in this area that rely heavily on manual intervention resulting in delayed intelligence. For example, a flight reconnaissance system where RDM embedded database technology is currently deployed. In this case, multiple vehicles are updated with their mission data via a ground system by uploading the data into a RDM embedded database on a removable storage device (hard drive). Once loaded, a soldier takes the device and physically walks it from the ground system to the vehicle prior to the start of the mission. In addition, this same RDM embedded database is used throughout the mission to gather flight data and events. At the completion of the mission the drive is again physically removed from the vehicle and walked back to the ground stations, where the data is aggregated and analyzed.

There are a couple of requirements that drive the decisions for the current solution. First, security is absolutely maintained by not electronically transferring the information. Secondly, implementation costs are kept to a minimum by placing as little code onto the airborne vehicle thus reducing the required certification and significant costs associated with it.

RDM Embedded can be configured as an extremely small footprint database engine keeping the certification costs to a minimum without sacrificing the functionality required to achieve secure data communication, system reliability and recovery through ACID transactions, and efficient and flexible data management based on the network data model. Combine this with real-time dataflow and the airborne systems can receive its flight data electronically prior to takeoff and any reconnaissance data can be passed back to the ground station while the vehicle is in-flight for real-time processing. With real-time intelligence real-time mission changes can take place without delays.

The Live Building Power Monitoring Demo

Please contact sales@raima.com for a live demonstration of the dataFlow technology. The demonstration is illustrated below using three individual single board computers, a netbook, a desktop computer, and a network for intra-computer communication.



Each embedded computer consists of a few separate components: RDM Embedded database, microHTTP Server, the replication server, and an application. Utilizing RDM Embedded's unique network model the application objects managed within the database represent buildings, units and measurements organized as hierarchies. Every second the application randomly generates a new power measurement for each of the units defined in the database. The application also makes sure to circulate the measurements so that each unit has a maximum of 20 random entries at any given time.

Whenever a new data point is added a RDM Embedded transaction is created which is instantly picked up by the replication server and handed off to the replication client. At any time during this process a web browser can surf the microHTTP server to add or delete building and unit data. The server also exposes all of the measurement data that is displayed through the browser using AJAX calls to populate a chart object.

On the RDM Server side, all transactions picked up by the replication clients are translated into SQL insert statements and aggregated into the RDM Server database using the native SAG CLI API. The RDM Server database can be accessed by any SQL tool, but for our purpose we expose the data in near real-time using PHP and the Apache WEB server. All insert transactions created by processing data in the replication clients are also picked up by a transaction plug-in, which in our implementation, translates them into ODBC calls to the instance of Microsoft® SQL Server® 2005 Express server. These calls insert the data into the SQL Server database. Finally, on the Microsoft® SQL Server® side of the system, an internet information server is hosted exposing the data through ASP.Net.



The above screenshot shows how the 5 different web servers are aggregated into a single view.

Conclusion

As system complexity and cost pressure continues to increase, relying on third party software vendors with a proven history of solving the most harsh data management problems may be beneficial. This document has use cases where secure and reliable data management, data flow, and distribution are core requirements. There is enough said about return on investment, time-to-market, and quality by selecting commercial vendors for your project on the internet, and this is no exception. The data management and dataflow tasks outlined in this document are far from easy to tackle. Transaction management, dataflow, data distribution, data reliability and real-time data access have been Raima's primary business the last decades.

Contact Information

Website: <http://www.raima.com>

WORLDWIDE

Raima, a Birdstep Company
2101 Fourth Avenue
Suite 240
Seattle, WA 98121
Tele: +1 206 748 5300
Fax: +1 206 748 5200
E-mail: sales@raima.com

EUROPE

Raima, a Birdstep Company
Fridtjof Nansens Plass 5
0160 Oslo, Norway
Tele: +47 90628326
Fax: +47 24134701
E-mail: sales@raima.com