Optimize Your Software Stack for Industrial Internet of Things Devices

According to Accenture, the Industrial Internet of Things (IIoT) will help companies boost revenues by increasing production, driving innovation and transforming their workforce.[1] The most conservative estimates place spending on the IIoT worldwide at $500 billion by 2020. GE calculated the economic impact of just a 1 percent cost savings across industrial segments including aviation, power, healthcare, rail and oil & gas at $276 billion.[2] What are the ingredients that will help your team contribute to securing your company's piece of this enormous pie? And how can you ensure that you get your innovative products to market ahead of your competition while still offering your customers the best user experience?

Mentor Graphics, Raima, and Datalight have collaborated to define and measure data management for IIoT edge devices. This paper examines the benefits that come from optimizing the software components of the data storage stack in terms of performance, reliability and device lifetime.

Contents

Why Examine IIoT Software Stack?

Today we are seeing intelligent processing and data analysis being pushed out of the cloud, toward the edge. For example, data is now being stored and managed in controllers, gateways and all the way to sensors and controls in devices capturing data and responding to real-time events. In the past, data would be sent to the cloud, decisions would be made there, and data stored there. Now devices need massive reliable data storage so that the data analysis, event logging, alarm signaling, and diagnostic communication can happen in real time, on the device with only outcomes transmitted across the network.

"Any piece of equipment becomes a productivity-enhancing hybrid if it generates data and the operator puts that data to work as a digital service within its supply chain."

Accenture,
"Driving Unconventional Growth Through IIoT"

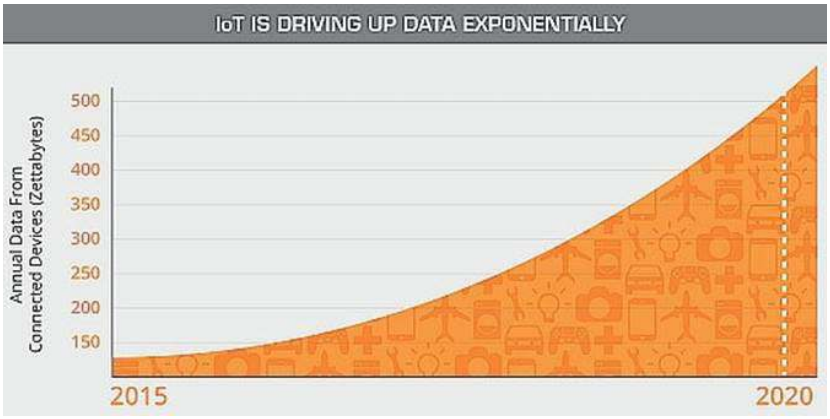

Figure 1:Source EDN

Production of data on devices at the edge of the IIoT is expected to exceed 500 zettabytes annually by 2020.[3] Jet engines generate as much as 10 GB/s of potentially useful diagnostic and performance-enhancing data.[4] A single residential smart-meter on the power grid is capable of managing upwards of 20 MB of data per hour, using it to optimize efficiency and reduce the environmental impact of energy production.

In a recent report, titled *Driving Unconventional Growth through IIoT*[3], strategy consulting firm Accenture made the case for IIoT devices opening new opportunities for companies throughout the ecosystem and remarked that "any piece of equipment becomes a productivity-enhancing hybrid if it generates data and the operator puts that data to work as a digital service within its supply chain."

We couldn't agree more. Data is becoming a key enabler in industries as diverse as energy, medicine, agriculture and transportation. And as more and more devices rely on data captured by sensors, and make decisions based upon it, the quality of that data becomes mission critical.

Get to Market Quickly with Pre-Tested Configurations

To capitalize on the opportunity presented by the IIoT, developers of devices need to get their projects into the marketplace quickly, with the functionality and performance that will delight their customers for years to come.

Embedded Linux is the leading choice for modern developers, offering designers of IIoT devices a rich landscape of parts and pieces to build virtually anything. Finding the best combination of open source components for your project can be time-consuming, and compatibility issues – or bugs – can derail product delivery. Mentor® Embedded Linux® is a pre-tested, commercially supported, extensible, and customizable platform ideally suited for IIoT devices that enables teams to quickly and efficiently scale to productivity.

Leverage an Embedded Database for Decision-Making Efficiency

The processors and systems running today's IIoT devices are more powerful than the mainframes on which databases were first implemented. Linux is a more sophisticated operating system and is a powerful platform for database systems. Linux on an embedded computer can support the capture, storage, analysis and transfer of huge volumes of data.

Most IIoT devices monitor and control live conditions, such as engine temperature, vibration and fuel-consumption. They must respond to extraordinary circumstances, and the best responses are not simply algorithmic — they are based on a history of observations. For example, many machines have sensors that capture temperatures or vibrations as indicators of the machine's condition and its need for maintenance. An embedded computer can have a database of waveforms and limits that raise maintenance signals, or even better, suggest different operational parameters that avoid the need for maintenance or extend the life of the machine.

Database management is the service that models complex data, optimizes the storage on permanent media, performs analysis and facilitates decision making on the computing device that is nearest to the data and controls over a physical system.

Don't Compromise Reliability

Many designers choose a database to make organizing and accessing data easier and more power failsafe. However, the methods used to ensure reliability can cause premature failure

of your storage devices and may also impede performance. In this study we examine methods for optimizing performance and ensuring a long and happy life for your storage media.

A defining difference between today's embedded computer and traditional mainframe is administrative support. An embedded device in an IIoT environment may be expected to perform without intervention or maintenance for many years, if not decades. A Database Management System (DBMS) in the embedded device must not interfere with ongoing instant-response operation, fill the storage media, consume computer memory, or contribute to the wear and tear of the computer components, especially flash memory-based storage.

The economic impact of unplanned downtime, which can be exacerbated by data loss or corruption, has been estimated at $45 million per day in aviation, $7 million per week in oil & gas, $400 million per day for class 1 railroad, and $1 million per day in lost revenue per refinery. Reliable data storage improves the accuracy of asset tracking which can lessen these costs by triggering preventive maintenance and reducing downtime.

Databases are designed to save data safely so that even in rough and unreliable power conditions or on-and-off devices, the saved data is always there. The set of properties that define database reliability can be described by the acronym ACID (Atomicity, Consistency, Isolation, Durability.) The steps that must be executed by a database to achieve this using most filesystems often results in multiple, redundant writes, which shortens the lifetime of the storage media. Raima Database Manager (RDM) combined with the Datalight Reliance Nitro file system is able to save its data safely and atomically while minimizing the frequency and size of flash device writes.

Reliance Nitro is a transactional, copy-on-write file system and gives developers (or applications like RDM) full control over when data is committed to the media. When Reliance Nitro performs a transaction it writes the data to a fresh area of the media and does so in a fully atomic fashion. If a transaction does not complete, for instance in the event of a power failure, the data from the previously committed transaction is intact. In this way, Reliance Nitro provides atomicity, durability and keeps the metadata and file data consistent.

By relying on the surety provided by Reliance Nitro, RDM can reduce the redundancy of data that is written.

Think Beyond Local Data on the Device

RDM is able to run in a disconnected, standalone environment as a complete, functional DBMS workhorse on IIoT devices. In a connected environment, RDM becomes part of a "data grid," able to receive, process, store and pass on data to other computers, devices, or cloud-based servers, altogether solving very big problems.

This is because RDM can use the built-in local data store or the data store on other computers. Its remote login feature allows databases to be shared among multiple computers. In addition, RDM is able to replicate local transactions to remote databases (RDM or otherwise) in an intermittently-connected environment, pushing changes out to the grid or cloud whenever the connection is present.

These are the low-level functional parts out of which a cloud solution is built, where one or more cloud-based databases are accessed only by name regardless of their actual location.

Optimize for Performance

In this study we used a benchmarking program, TPC-B, programmed by Raima to test the performance of RDM in various environments. TPC-B is a standardized benchmark approved by the Transaction Processing Performance Council. It has always been characterized as a "stress test" for databases because it includes:

- Significant disk input and output
- Moderate system and application execution time
- Transaction integrity

TPC-B is especially appropriate for measuring the performance of an embedded computer because it has no concept of "users." Instead, it receives a steady stream of concurrent transactions, saving the data and verifying its integrity as it goes. Results are expressed as transactions per second.

Tests were run on an NXP i.MX6 development board running Mentor Embedded Linux, configured to use Datalight Reliance Nitro or ext4 file system in separate test runs.

Figure 2: NXP i.MX6 reference platform used for benchmarking

RDM performs significantly faster when utilizing a transactional file system like Reliance Nitro. With additional throughput available, more data points can be captured, potentially increasing the prediction capabilities and certainly providing a more responsive user experience.

These results are exciting in two ways. First, we found that the Reliance Nitro file system performed significantly faster even when its transactional functionality was not being exploited. In other words, Reliance Nitro ran faster as a regular file system than the standard Linux ext4 file system. But when Raima altered RDM to exploit the Reliance Nitro transactional functionality, the numbers improved even more.

RDM performs transactions safely on standard file systems, incurring overhead because it cannot assume that everything written to storage will be present if power is lost. There are times when RDM is forced to wait until it has confirmation that the data is safe - this is the essence of transaction processing. Reliance Nitro moves the transaction processing into the file system itself where

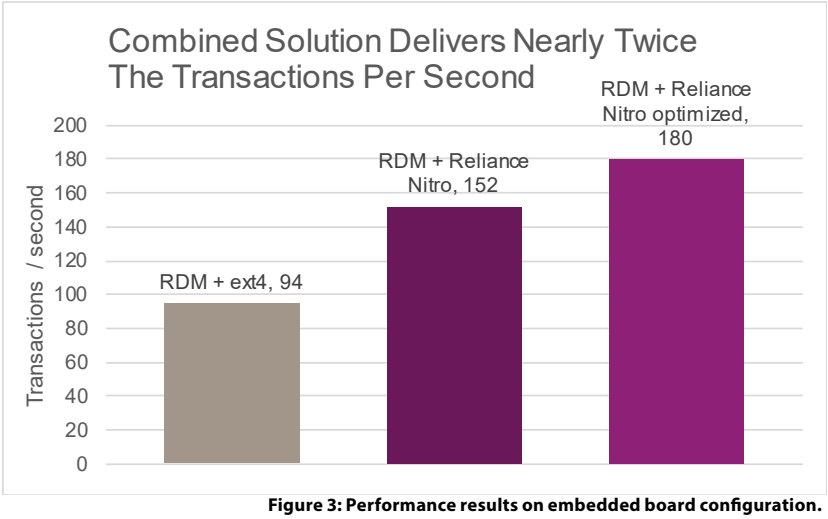

Figure 3: Performance results on embedded board configuration. Larger numbers are better.

it can commit a transaction more efficiently than RDM. So when RDM defers to Reliance Nitro to keep the data safe, it is able to do much less work. The result is a benchmark that runs twice as fast with the same reliability!

Protect Lifetime of Vulnerable Flash Memory-based Storage

Most IIoT devices will use some form of flash memory-based storage: SSD, SD, eMMC, eUSB, NAND, or NOR. These components are vulnerable to premature failure if not used efficiently. The benefits of using RDM and Reliance Nitro together doesn't stop with performance and reliability but can extend the lifetime of storage as well.

Using the embedded configuration, we measured total erases (the key indicator of flash media life expectancy) that occurred during each test run, and we calculated the average of transactions completed during each erase.

With the default configuration (making no changes to RDM), Reliance Nitro showed erase counts of 21 percent lower than ext4, while performing 60 percent more database transactions, doubling the transactions per erase.

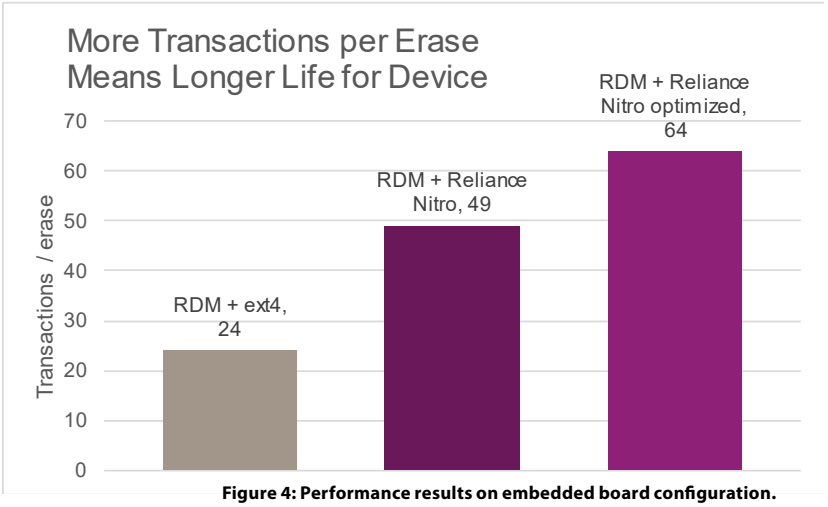

Figure 4: Performance results on embedded board configuration. Larger numbers are better.

When database behavior was optimized to take advantage of the Reliance Nitro transaction model by relying on a single file-system-wide commit rather than the individual file commits RDM must perform on non-transactional file systems, the system executed twice the database transactions with 75 percent of the erases of ext4, delivering close to three times the transactions per erase. Making these same alterations when using another file system would result in data corruption.

Conclusion

Best-in-class IIoT devices offer superior performance, long lifespan, data integrity, and get to market ahead of their competition. By understanding how the components in your software stack work together and reducing redundant operations, you can improve system efficiency and performance without sacrificing reliability.

Most file systems do not comply with ACID properties and therefore databases must take steps to ensure the Atomicity, Consistency, Isolation and Durability of transactions. This often results in redundent writes which can lead to premature media failure. Our testing shows that using a transactional, copy-on-write file system ensures ACID compliance without redundant writes.

Ultimately, the customer benefits with an optimized software stack that offers a quick-to-market solution, improved warranty costs, and better business outcomes.

About the Contributing Companies and Products:

Datalight, Inc

When data integrity, time to market and design flexibility matter, the world's leading device manufacturers choose Datalight. Datalight's software ensures that IIoT devices have power failsafe reliability throughout their entire lifetime, reducing downtime and warranty (and out-of-warranty) costs. Datalight products have shipped in hundreds of millions of embedded devices in demanding product categories, such as automotive, medical, retail, industrial automation and military/aerospace.

Our product line includes:

- Reliance Nitro, a transactional power failsafe file system for sophisticated embedded systems
- Reliance Edge, a tiny, power failsafe file system for microcontroller-based devices
- FlashFX Tera, comprehensive software management for raw flash.

Visit www.datalight.com/products/embedded-file-systems/reliance-nitro for more about Reliance Nitro

Mentor Graphics Corporation

The Mentor Graphics Embedded Systems Division enables embedded development for a variety of applications including automotive, industrial, smart energy, medical devices, and consumer electronics. Embedded developers can create systems with the latest processors and micro-controllers with commercially supported and customizable Linux-based solutions including the industry-leading Sourcery CodeBench and Mentor Embedded Linux products. For real-time systems, developers can take advantage of the small-foot-print and low-power-capable Nucleus RTOS. For more information, visit www.mentor.com/embedded

- Mentor® Embedded Linux® is a commercially supported, extensible , and customizable platform ideally suited for Industrial Internet of Things, that enables teams to quickly and efficiently scale to productivity
- Mentor Embedded Linux provides connectivity for IIoT that enables the data to get on the device from the enterprise down, through controllers and other connected devices in the Industrial environment. It also enables you to secure the connectivity from the enterprise to the controller layers down to the edge devices.

Raima, Inc

Raima provides tools and expertise to coordinate data from the wide variety of sensors and other inputs, simplifying access. Learn more about Raima Database Manager at www.Raima.com

[1] Accenture Techonology (2015). *Driving Unconventional Growth Through the Industrial Internet of Things*. Retrieved from Accenture website: https://www.accenture.com/us-en/_acnmedia/Accenture/next-gen/reassembling-industry/pdf/Accenture-Driving-Unconventional-Growth-through-IIoT.pdf

[2] GE. (2016, June 10). *[No Unplanned Downtime] [Infographic]*. Retrieved from website: https://www.predix.com/ge-industrial-internet-infographic?utm_source=II_Survey_PDF&utm_medium=PDF&utm_term=inbound%2520PDF%2520clicks&utm_content=None&utm_campaign=II_Survey_PDF

[3] Architects of Modern Power (2016). *[Drivers of Digital Power] [infographic]*. Retrieved from website: http://www.ampgroup.com/the-drivers-of-digital-power/

[4] VR World (2015, May 8). *BIG DATA IN PLANES: NEW P&W GTF ENGINE TELEMETRY TO GENERATE 10GB/S*. Retrieved from website: http://vrworld.com/2015/05/08/big-data-in-planes-new-pw-gtf-engine-telemetry-to-generate-10gbs/

Datalight

Optimize Your Software Stack for Industrial Internet of Things Devices

by
Kerri McConnell, Vice President, Sales & Marketing, Datalight
Wayne Warren, Chief Technology Officer, Raima
Kathleen Tufto, Senior Product Manager, Mentor Graphics

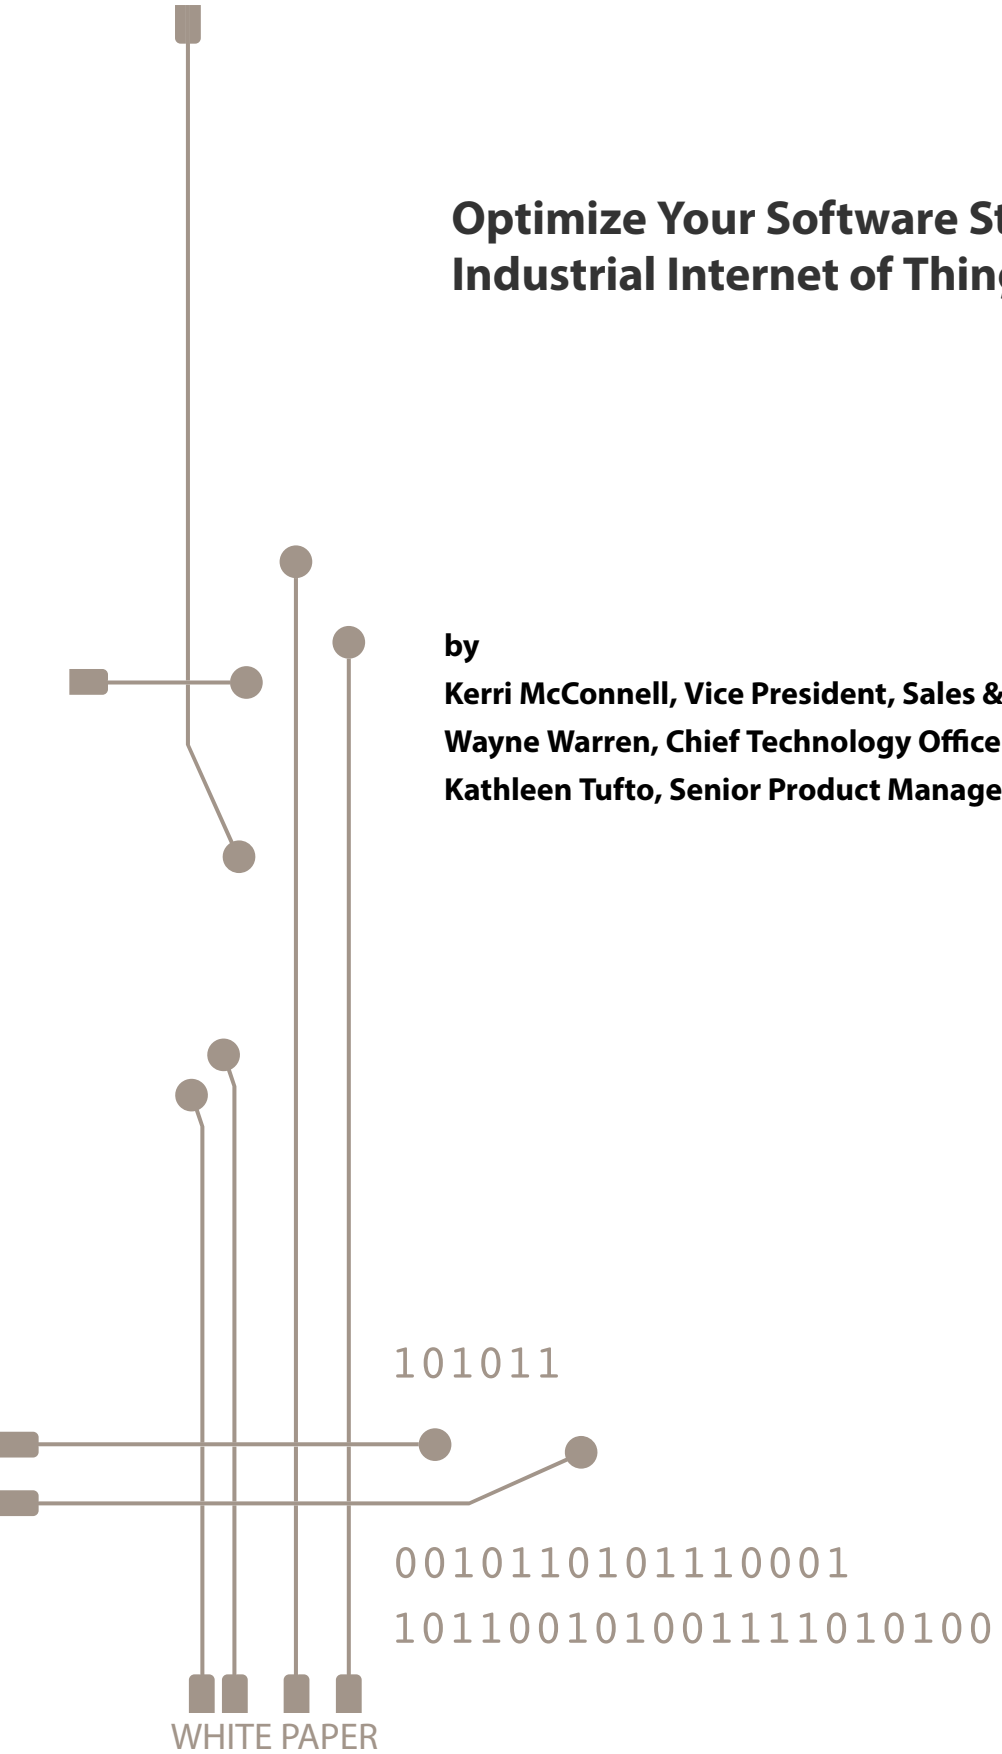

WHITE PAPER