Sun and Auto—ID

Enabling Smart Objects: Breakthrough RFID-enabled Supply Chain Execution Infrastructure

Enabling Smart Objects: Breakthrough RFID-enabled Supply Chain Execution Infrastructure

"Our solution strategy starts with R&D that is production-proven in our iForce centers and ends with mass customization via Sun Services and our iForce partner community. We are systematically accelerating the time to service and ROI by capturing the most frequent design patterns and methodologies for a "step and repeat" effect in our customer base. Unlike our competitors, we are putting structure into our Infrastructure Solutions."

–**Scott McNealy**
Chairman, President & CEO
Sun Microsystems, Inc.

Table of Contents

Chapter 1

Executive Summary

The emerging set of global standards and specifications for automatic identification can bring unprecedented visibility in the supply and distribution chain. Using technology breakthroughs in radio frequency identification (RFID) design, the Massachusetts Institute of Technology (MIT) Auto-ID Center, along with the Uniform Code Council (UCC), is leading a group of more than 90 companies and research centers to define widely supported global standards in reading, finding, and formatting product information. These standards are being designed for use as a next generation of the bar code. The Auto-ID standards will create a cost-effective way to make the supply chain more efficient. The compelling aspect of an Auto-ID enabled operation is the association of information with product movement. The combination of tags, antennas, readers, and local computers ("Savants") provides a near real-time view of product status and location. Many companies have begun trials to determine how this new infrastructure can be best used to make significant improvements in enterprise cost structures or revenue capabilities.

RFID technology has been around for many decades. The Auto-ID initiative addresses two areas: lowering the cost of the tags and infrastructure needed to read them, and standardizing interfaces to enable interoperability and multi-vendor implementations. This will broaden the adoption of RFID by making it cost effective to apply tags to even low cost items, like soda cans or paper towels, and enabling companies to share RFID infrastructure. Before Auto-ID, each implementation was based on proprietary technology provided by the solution developer. The Auto-ID standards will help ensure that different companies in a supply chain ecosystem don't need to use the same solution developer to have their RFID equipment interoperate. This equipment will enable companies to send information along with products. This information will need to be integrated into the enterprise information system (EIS), enterprise resource planning (ERP), supply chain management (SCM), warehouse management (WM), customer relationship management (CRM), and other similar applications.

This paper provides an overview of the Auto-ID specification and its components, what Sun has learned from initial trials, and what products and technologies can be used to create an Auto-ID implementation. Taking advantage of more than five years of research and development investment and proactive collaboration with Auto-ID and its key members, Sun Microsystems and its iForce℠ Partners deliver a breakthrough RFID-enabled supply chain execution infrastructure that will help global enterprises dramatically improve the visibility, velocity, and value of their supply chains.

Intended Audience

This document gives a technical description of the Auto-ID infrastructure and is aimed at IT professionals who are considering an Auto-ID implementation. However the Auto-ID overview and deployment considerations are relevant to the line of business professional that wants to become familiar with the infrastructure needed to deploy a solution.

"Gillette expects that Auto-ID will help in reducing shrinkage (product losses) and improving retailer relations. We expect the current impact on these to be approximately three to five percent of annual revenues."
– Dick Cantwell, *VP of Supply Chain, Gillette*

"Proctor & Gamble expects to save $1.5 billion annually in supply chain costs through the use of Auto-ID."
– Steve David, CIO, *Procter & Gamble*

Chapter 2

Auto-ID Overview

Automatic identification (Auto-ID), is a set of industry-supported specifications that are being developed for RFID systems. These standards cover the basic technologies and protocols that enable IT systems to identify and track virtually any physical thing. The initiative, led by the Auto-ID Center at the Massachusetts Institute of Technology (MIT), is also supported by leading companies in many industries. Founded in 1999, the Auto-ID Center was created in response to the Universal Code Council's vision to build the next-generation bar code. The Auto-ID Center's mission is to design an infrastructure and develop standards for a universal, open network that can identify individual products and track them through the global supply chain. The long-term goal is a universally adopted standard that will be used by companies to associate information with trillions of objects such as retail products, manufacturing materials, medicine, and much more.

Fundamentally, Auto-ID is about tagging objects so that they may be identified—ideally without human intervention—and reading each tag in order to identify and forward information about each item. This data can then be used to automatically identify objects, monitor location and movement, and trigger events.

The Auto-ID standards are expected to be announced in September 2003. At that time, they will be rebranded as the EPC-System. In this document we continue to refer to it as Auto-ID.

By offering companies an unprecedented real-time view of inventory throughout the supply chain, Auto-ID is expected to enable significant gains in operational efficiencies. For example, retailers will have nearly nearly instantaneous, rather than monthly or quarterly, data about which products have been sold. This information could be used to more tightly control inventories and thereby reduce costs.

Chapter 3

Architecture and Technology

The Auto-ID Center and its sponsors are working to design a standard that will enable interoperability of solutions from multiple vendors and bring component cost down to a level where the new technology can be widely deployed. The key components of the Auto-ID standard are:

- Electronic Product Code (EPC)
- Radio frequency identification (RFID) tags
- Tag readers
- Savant servers
- Object Name Service (ONS)
- Physical Markup Language (PML)

Electronic Product Code (EPC)

A numbering scheme called *electronic product code* (EPC) will allow assignment of a unique identifier for any physical object in the world. Unlike the bar code that identifies product categories (for example, a 330-ml can of cola), the EPC is capable of uniquely identifying every single item that rolls off a manufacturing line. One example for how this could be used is to easily and quickly identify which products might be nearing their expiration dates. The EPC standard will provide a migration path for systems using bar codes

01	• 0000A89	• 00016F	• 000169DC0
Header 8-bits	**EPC Manager 28-bits**	**Object Class 24-bits**	**Serial Number 36-bits**

Figure 1: Electronic Product Code Fields in EPC Type I Format 96 bit

The EPC includes a header and three data fields, as shown in Figure 1.

- The header identifies the EPC's version number, enabling different types and future revisions of EPC.

- EPC Manager is the name of the enterprise—typically the manufacturer's name.
- The object class refers to the class of product, usually the stock keeping unit (SKU). For example, "Arrowhead Water, 500-ml plastic bottle, California Redemption version."
- The serial number is unique to the item instance. For example, it could identify a specific brand, size, and bottle of drinking water.

The Auto-ID Center has proposed both 64- and 96-bit EPCs. The 96-bit EPC provides unique identifiers for 268 million companies. Each company can have 16 million object classes, with 68 billion serial numbers in each class. Since this many serial numbers are not needed at this time, a 64-bit code has been proposed to help keep down the price of RFID chips. To address future requirements, 256-bit codes are also specified.

Radio Frequency Identification Tags (RFID)

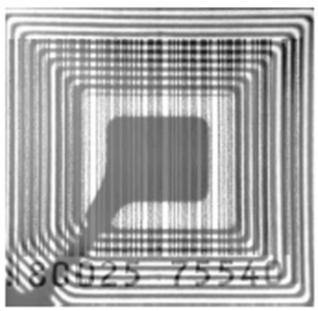

Figure 2: RFID Tags Include Chip (center) and Antenna (surrounding)

A tag is made up of both an RFID chip and an antenna. Three different types of tags—active, semi-passive, and passive—are used in different applications. See Table 1 below.

	Active Tag	Semi-passive Tag	Passive Tag
Power Source	Battery on tag.	Battery for chip operation. Radio wave energy from Reader for communication.	Radio wave energy from Reader for operation and communication.
Tag Signal Availability	Always on, 100 feet	Only within field of reader	Only within field of reader, less than 10 feet
Signal Strength Tag	High	Low	Very low
Required Signal Strength from Reader	Very low	Low	Very high
Typical Applications	Useful for tracking high-value goods that need to be scanned over long ranges. Example: railway cars on a track.		Useful for high-volume goods, where items can be read from short ranges. Example: retail check out.

Table 1—RFID Tags and Characteristics

Much of the attention surrounding Auto-ID has been focused on the radio frequency identification (RFID) tag technology. To gain widespread adoption, tags must be reduced in price. The last year has already seen tag prices drop from fifty cents to less than ten cents each. Already, RFID chips can be manufactured as small as 0.3 millimeters —approximately the size of a pencil point. The Auto-ID Center is focussing on passive tags—though other tag technologies are not excluded from the specifications—as they are expected to be far less expensive to make and use.

In concept, RFID tags have an advantage over bar codes. Wireless RFID tags can be read anywhere within range of an antenna, while bar codes must be read by a scanner in the line of sight.

Chips in RFID tags may be read-only or read-write, depending on application requirements. Some tags employ EEPROM technology, which enables data to be written to the tags using a special electronic process. Companies can write to an EPC tag when the item is produced and packaged; however, for all practical purposes the tag is a read-only after it has been manufactured.

Due to varying government regulations, there are regional operating frequency differences for RFID systems. Geographic areas, such as Europe, North America, and Asia, have specified different operating frequencies for the EPC tags to operate. For example, tags designed for the North American market typically operate at 915 MHz. European tags operate at different frequencies, including 13.56 MHz, 800-1000 MHz (UHF) and 2.4 GHz frequency bands.

Tag Readers

RFID readers use a variety of methods to communicate with tags. The most common way to read passive tags at close range is called inductive coupling (the same technology used for key card entry at many companies). The coiled antenna of the reader creates a magnetic field with the coiled antenna of the tag. The tag draws enough energy from this field to send back its EPC. Readers cost about $1,000 today, but are expected to be available for $100 in volume later on. Readers are available which operate with tags designed for different regions.

The Savant

In an environment where readers are picking up a nearly constant stream of EPCs, managing the data is a significant task. The Auto-ID Center has designed software technology called The Savant to act as the leading edge of this infrastructure. The Savant distributed architecture gathers, stores, and processes EPC data from one or more readers. Savants also smooth data, correct duplicate reader entries, intelligently store and forward data up or down the chain, and monitor for events (low stock level, for example). Then, Savant servers pass data up to the ERP systems, through a full time connection, or synchronizing data on an "as needed" basis. See Figure 3 for an example of where Savant technology fits in the Auto-ID solution architecture.

The Savant framework consists of a set of geographically distributed servers. Savants are connected to RFID readers which continuously collect EPC data from tags, and feed this data to the Savant. For each reading, the Savant maintains information such as:

- EPC of the tag read
- EPC of the reader (reader EPC) that scanned the tag
- Time stamp of the reading
- Non EPC-related information, such as temperature or geographical position, that is observed by the reader

The Savant consists of three major software modules:

1. **Event Management System (EMS):** A Java™ technology-based system provides an API for various types of readers, collects data in a standard format, allows customized filters to smooth and clean data, and provides various mechanisms to log data into a database or remote servers using standard protocols (HTTP, SOAP, Java™ Message Service, and Java Message Queue).
2. **Real-time, in-memory data structure (RIED):** The RIED can be used to store event information by Edge Savants. RIED provides the same interface as a database, but offers much better performance. Applications can access RIED using JDBC™ technology or a native Java technology interface. RIED also supports SQL operations and can maintain "snapshots" of the database at different time stamps.
3. **Task Management System (TMS):** The TMS provides an external interface to schedule tasks. It simplifies the maintenance of distributed Savants, because the enterprise can maintain Savants by merely keeping the tasks on a set of class servers up to date, and appropriately scheduling tasks on the Savants. In addition to data gathering and transmission, the TMS can be used to request PML and ONS activity, and schedule and administer tasks on other Savants.

Note that at the time this document was created, there is an effort to specify just the payloads and protocols for connecting to a Savant, rather than specifying internal software functionality.

Software Components

Object Name Service (ONS)

ONS matches the EPC code to information about the product. Once a Savant receives EPC data, it can query an ONS server to find out where product information is stored. This system has been modeled after the highly scalable, reliable, and high-performance Domain Name System (DNS) system used in the Internet. When using the Internet, one can provide a URL to a DNS server and expect to receive the associated IP address in less than a second. Because it is anticipated that Auto-ID will scale to identifying trillions of products from thousands of companies using public network infrastructure, DNS design concepts were used in the design of ONS.

Physical Markup Language (PML)

The EPC identifies individual products, but useful information about the product is written in a new, standard computer language called Physical Markup Language (PML). PML is based on the widely accepted, extensible markup language (XML), and is expected to become a universal standard for describing physical objects, processes, and environments. Thus PML can store any information that could be useful; for example, product composition, lot number, and date of manufacture. This information can be used to create new services and strategies. For example, a consumer could find out how to recycle a product's packaging, a retailer could set a trigger to lower prices on milk as expiration dates approach, or a manufacturer could recall a specific lot of product. PML is designed to be a dynamic data structure, with information that can be updated over time. For example, the PML record for a product can be updated to store the location of a product as it moves through a supply chain.

Putting It All Together — The Auto-ID Architecture

The RFID-enabled Supply Chain Execution Infrastructure co-developed with Auto-ID members is designed to be highly flexible — open to different implementations, technologies, and protocols. The Sun™ solution architecture illustrated in Figure 3, highlights a very flexible design based on open standards. Once EPC data are detected by the readers, they are passed to The Savant. The Savant acts as event manager, filtering out extraneous EPC reads or events. The ONS Server provides the IP address of a PML Server that stores information pertinent to the EPC.

Data from the Savant is passed into the application infrastructure, or operations bus, either locally or over a WAN such as the Internet. From here, the data is made available to virtually any application that can make use of it.

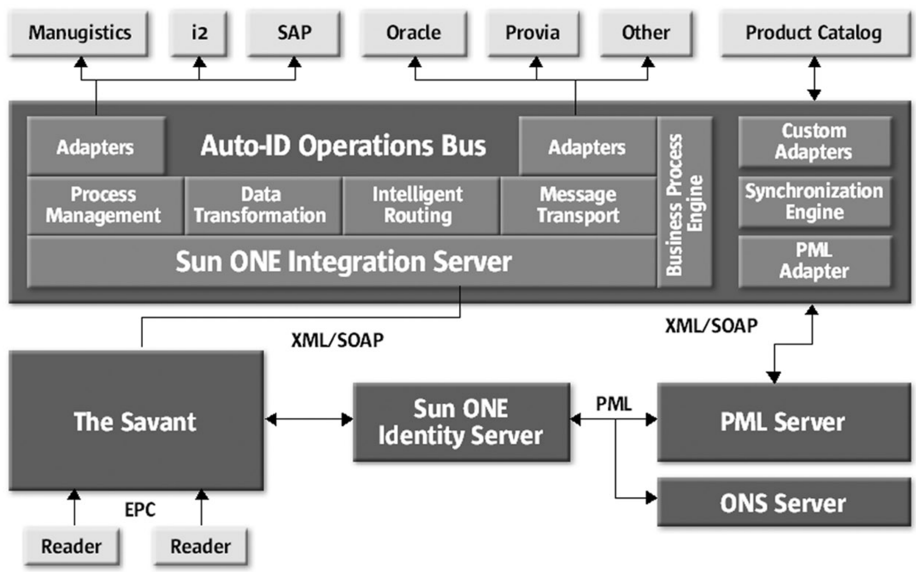

Figure 3: Sun Auto-ID Solution Architecture

Chapter 4

IT and Deployment Considerations

Even though this is a new market, many leading companies have shown that they are willing to invest in Auto-ID technology. Sun has been actively engaged with multiple large manufacturing and retailing enterprises in their efforts to integrate Auto-ID technology into supply chain IT systems. A number of pilots are underway, and it has been widely reported that Gillette agreed to purchase 500 million RFID tags. With more than 90 sponsoring companies—representing combined annual revenues of over $1.5 trillion and hundreds of billions of goods in the supply chain—it is clear that the use of this technology offers important business benefits and significant cost savings.

Some of these early efforts have been in the area of consumer packaged goods, including both retailers and manufacturers. Feedback from these companies indicates that they want to deploy the technology first in areas where they can realize a return on investment by improving internal operations, such as manufacturer-to-retailer shipping operations. They expect these gains to come from reducing out-of-stock goods and cutting labor costs as well as from reducing shrinkage and inventories.

Much has been learned from these early efforts. Key observations are outlined below.

Understand the Business Case

Before starting any project, a clear understanding of the use cases where Auto-ID will be applied is required. This should be followed by a thorough inventory of the operations and business processes,that will be impacted. Next, it is necessary to know what data is sent into enterprise applications, such as warehouse management systems, supply chain management (SCM), and enterprise resource planning (ERP) applications. Once the business elements are understood, the technology for an effective EIS solution can be properly applied. Some of these projects are being implemented alongside existing systems for tracking product movement (such as barcode systems). This strategy provides a well defined business process to be replicated with a new technology. This avoids having to reengineer a process, and the integrity of the new system can be verified by correlating results of the two systems.

Start With a Well-Defined Project

While some envision a time when Auto-ID will provide end-to-end visibility across the entire channel, enterprises are selecting projects that are limited in scope for test trials. Some manufacturers are using warehouse product tracking as the area to explore. Some retailers are exploring real time inventory tracking. In these situations, businesses must determine what event, or series of events, needs to occur before sending data to the EIS. For example, in an inventory application, data on when a product arrives and departs as well as its location are important. In a warehouse management application, matching a product to an order number is key. One consumer product manufacturer expects to start by tagging pallets, then cartons, before moving on to individual products. These early adopters are applying the lessons they learn in the early phases to other deployments.

Understand and Explore the Physics

Auto-ID solutions involve understanding the deployment of radio frequency transmitters, antennas, and tags. Tag technology selection is affected by the material composition, product shape, and packaging. Placement of tags on packaging and pallets will have a dramatic effect on reader performance. Environmental radio frequency noise must be understood and controlled in order to get accurate sensing of the tags. Antenna locations must be carefully chosen to optimize information gathering. These issues are collectively referred to as the physics of RFID. Expect that early explorations of the physics will be an important part of the solution implementation plan.

Plan Remote Maintenance

A thorough understanding of the deployment environment is essential. While many deployments will be different—warehouse, retail outlet, manufacturing floor, to name a few—a common theme is the lack of dedicated IT support at every location. It is not reasonable to expect untrained workers, whether they drive a forklift or operate a cash register, to maintain or monitor the system. Remotely managed, auto-discovering, self-healing, and self-provisioning technology is essential for successful deployments.

Don't Overwhelm Your Network

Early trials have shown that EPC readers can generate significant amounts of data. In one early trial, engineers connected readers directly to the EIS application, bypassing the Savants. The data overwhelmed the system, crashing it within minutes and taking other applications down with it. Savants are intended to buffer readers from the EIS infrastructure, filtering and processing EPC data as required. The Savants should be used to collect, filter, and upload the necessary data.

Savants are configured to unique applications with specific filters, reader adapters, business process rules, and so on. As with any IT component, Savants require periodic software updates and patches. Upon analysis, nearly 80 percent of the cost of one trial was due to technicians traveling to remote warehouses when updating and maintaining on-site equipment. The ability to deploy generic Savant hardware that can plug, play, and run with minimal on-site training and effort, is a key concept in controlling costs.

Perform Risk Analysis

Understanding the impact of these systems on operations is crucial. Businesses can lose tens of thousands of dollars per hour in revenue if systems at distribution centers and warehouses are offline. What happens to other systems if the project gets off to a suboptimal start? What are the service level requirements for the new system? Will the system become a business critical system upon which others rely? A deployment architecture will likely need to provide availability features such as redundancy and failover in a cost-effective design in order to meet business requirements.

Design in Flexibility

Cross-platform compatibility will be required in virtually any installation in order to work with existing application, platforms, and systems. A protocol- and schema-neutral design will enable the exploration of alternate implementations and work with the wide range of existing interfaces.

Auto-ID is expected to create more pressure to share data between enterprise trading partners. The electronic business extended markup language (ebXML) is an emerging framework for sharing data between business IT operations. ebXML provides loose coupling, is based on open standards, and results in applications that are easier to build, maintain, and reuse. A complete B2B framework, ebXML enables business collaboration through the sharing of Web-based business services. In addition to many other features and capabilities, ebXML provides a reliable messaging specification and facilitates declarative representation of business processes.

Implement Reliable Messaging

Reliable, interoperable message transport is also a key concept in moving business critical information across the network. The ability to send data exactly once is central to a reliable and trustworthy operation. The ebXML Messaging Service (ebMS) extends the SOAP specification to provide the security and reliability features required by production enterprise services. Instead of dedicated leased lines, businesses can use the Internet to reliably communicate between system components.

Interoperability with existing protocols will likely be a requirement. The Applicability Statement 2 (AS2) specification, which uses HTTP for transporting EDI transactions, is deployed in some supply chains. ebMS is a royalty free interoperable messaging specification that provides for exactly once quality of service guarantees. Work to extend ebMS is taking place as part of the WS-Reliability efforts.

Design for Security

As with any IT deployment, security is an important consideration. Data and control streams from many different locations may be coming into the data center through the Internet. Widely dispersed Auto-ID deployments may use the Internet as the primary means for connectivity. Remote management and monitoring of system components over the Internet must be implemented with an inherently secure architecture.

Security mechanisms such as encryption and authentication are critical to protecting enterprise data and overall system integrity. Enterprise security methods typically tend towards either a centralized or decentralized architecture (such as Kerberos and PKI, respectively). The Auto-ID Center does not specify a security technology, leaving that aspect of the implementation to the solution provider.

Privacy Considerations

There is a small but vocal group that has raised privacy concerns over the deployment of RFID solutions. Businesses that are planning for B2C trials and deployment should consider how this new capability will affect their corporate privacy policies. The most effective approach is to address these issues proactively at the beginning of any effort. A clear definition of how information will be collected must be created and consistently communicated.

The Auto-ID Operations Bus

The IT infrastructure at most enterprises is comprised of many different systems. Auto-ID is a new technology, and it is difficult to know all the places it may be productively employed in an enterprise IT architecture, or which EIS will want access to EPC data.

Early Auto-ID applications were deployed either as standalone applications, or the data feeds were connected on a one-to-one basis to existing EIS applications. Standalone applications create yet another set of reports, maintenance, support requirements, and so on. It is highly unlikely that an enterprise will want to develop new, separate applications that track product movement. It is much more likely that the existing EIS applications that need product movement information will use this as a new data source.

Connecting data feeds directly from Savants to each application can also be problematic. The direct connect approach can lead to information silos where product movement is controlled and only usable by a single application. If more than one application is directly connected to the Savants, this approach can result in the development and maintenance of multiple point to point links or connector "circuits." This will magnify software maintenance, upgrade, and replacement effort, increasing costs.

Rather than creating direct connections to specific applications, creating an operations data bus allows maximum flexibility when requirements, back end, or front end applications change, and avoids creating an information silo. Data feeds from the Savants are fed into the bus in a consistent manner, and any application that needs the data can easily access it using similar tools and techniques for each application. Flexibility and reliability can be increased, reaction time to business opportunity is reduced, and maintenance and support costs are reduced.

An additional benefit of this approach is that the operations bus layer can be used as a place to implement additional desired functionality like data warehousing, data reformatting or to apply business logic to events.

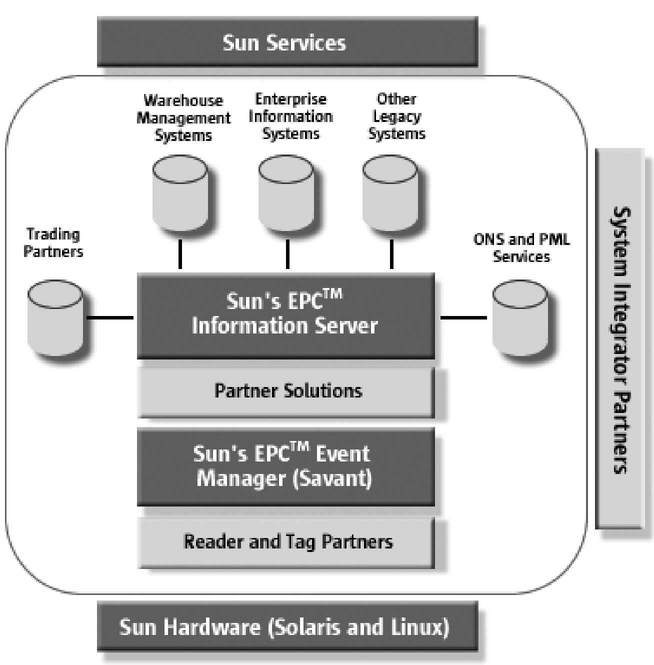

Figure 4: Using Sun's EPC Information Server as the Auto-ID Operations Bus

A pallet manufacturer—which generates revenue when its pallets are used for transit, not when they are stored at a customer site—implemented an RFID tracking system and collected the usage information in a data warehouse. Analysis showed interesting behavior patterns among its customers. Perishable goods customers typically returned the pallets promptly, while big-box retailers kept them stored in the parking lot for weeks or longer. Armed with this knowledge, the manufacturer changed its pricing model.

It may be required that Auto-ID events be stored in a data warehouse. This will allow later data mining and trend analysis. This can be a backup repository in case there is a need to recall raw product movement information. The operations bus approach provides a logical layer for implementing data warehousing functionality.

Information coming from the Savant layer may not be in a format that is understood by every back end application. The operations bus layer can be used to reformat or translate data into a form readable by the target application. This insulates the back end application from changes in data formats, and can avoid modifying the back end application to understand new data types.

It is possible that not all product movement events are needed by all applications. The operations bus can be a good place to make decisions about when to pass what information to back end applications. Business logic can be applied at this layer to process product movement information and thereby reduce network traffic and application load potentially created by passing extraneous information to the EISs.

Chapter 5

The Breakthrough Auto-ID Enabled Supply Chain Execution Infrastructure Solution

Enterprises have indicated that they are looking for reliable, available, scalable, manageable, and secure solutions that can be integrated into their existing IT infrastructure, and Auto-ID is expected to become a business-critical application. Sun has made a commitment to being a leader in existing and emerging technologies through its participation in standards committees. Sun was an early sponsor of the Auto-ID Center, and Sun employees currently head up The Technology Board and the Software Action Group. Sun provided technical, engineering, licensing and legal expertise to the Auto-ID Center and is actively involved in many pilots. Through this experience we have developed a proof of concept for Auto-ID in an enterprise. It will be described below.

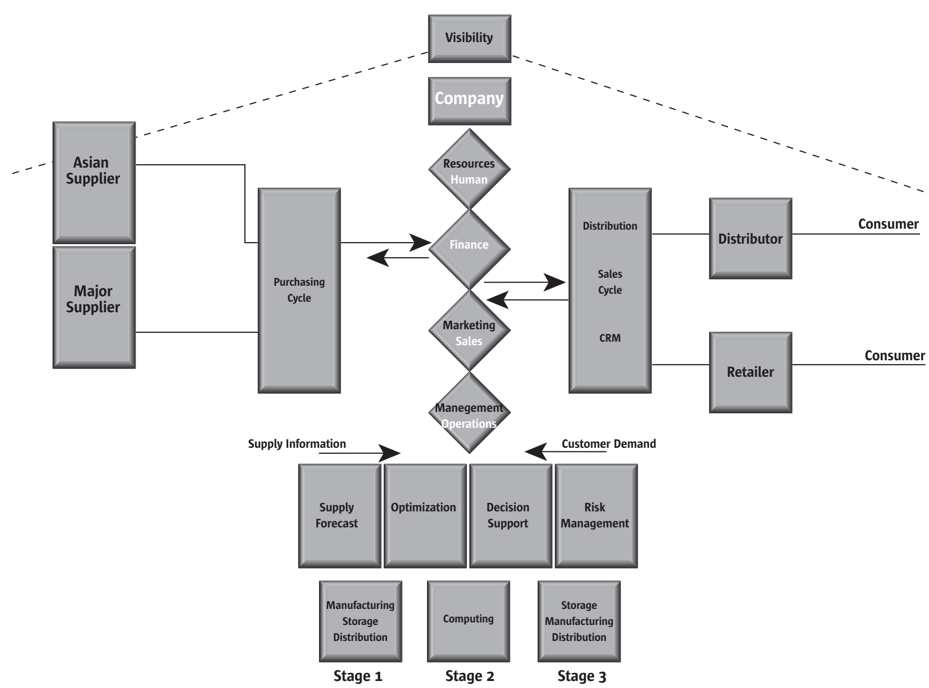

Figure 5: Sun and its iForce℠ Partners Enable Total Visibility Across the Supply Chain

The iForce Center Implementation—Real World Proof-Of-Concept

At the Sun iForce Center in Menlo Park, California, Sun and iForce Partners have built an Auto-ID system that implements the operation bus approach recommended earlier. The example modeled in this implementation is the Advanced Shipping Notice (ASN) process. It sends notification to the customer that the shipment has been sent, along with exact content details—a very accurate packing slip. Because the RFID tags can be read without unpacking a pallet, the contents of each pallet are instantly recorded as they pass through a loading dock door. The ASN application at the iForce Center is designed to accept information from many types of readers and Savants, over virtually any network. The data is received into an integration bus where it can be shared and leveraged among multiple back-end applications and repositories.

The ASN application is used to demonstrate how companies can use the Auto-ID infrastructure to automatically send accurate and detailed shipment information. Built on the Sun™ Open Net Environment (Sun™ ONE), it is based on open APIs and designed to be flexible, integrated, and integrateable. While the foundation software is built with Sun ONE software, (integrated), software from other vendors can be used (we've coined the term *integrateable* to describe this capability). In addition, software from iForce partners is also used.

Architectural Overview

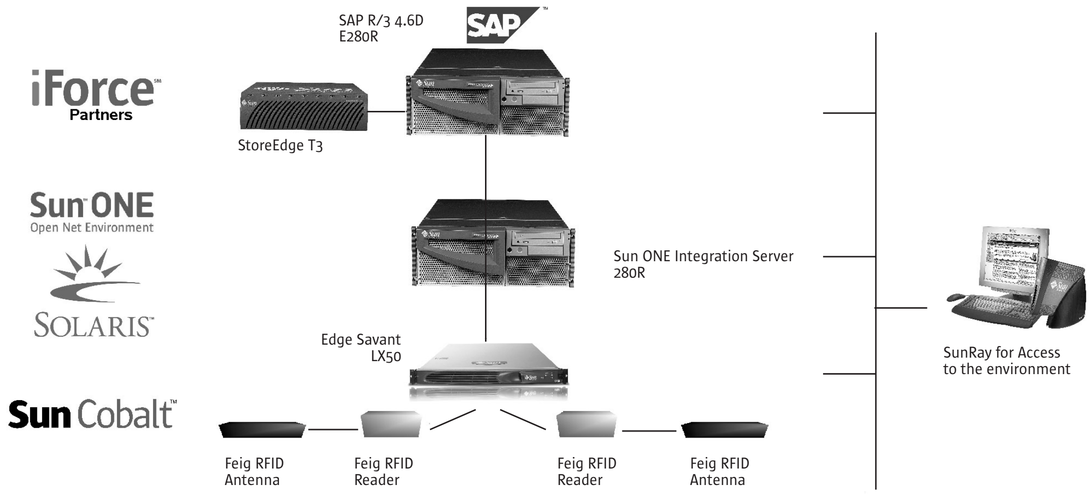

Figure 6: The Hardware Configuration of the iForce Center Implementation

At a high level, an RFID reader, Savant, operation/process layer, SAP order management applications, and inventory management system (IMS) application comprise the ASN prototype. The demonstration application shows how removing items from a shelf automatically triggers an order once an inventory threshold has been reached. The order is sent from the Savant to the process engine, which submits it to the SAP order system. When an order for product is entered into the SAP order management screen, it again interacts with the process engine, which relays a message to an IMS application in a simulated warehouse. At the warehouse, items from the order are placed in the staging area — which are accounted for against the order. Once the items match the order, the IMS notifies the process engine, which updates the SAP application. In the SAP order management screen, the order is marked as fulfilled, and an ASN is sent to the customer's ERP system.

The prototype highlights the advantages of an integration server at the operations layer — the horizontal architecture is flexible and integrateable, and works with many different back-end applications:

- The warehouse management system (WMS) application can communicate with other enterprise application environments, such as transportation and order management systems, without a significant amount of new code. Unlike creating multiple point-to-point links between the Savant and applications, and between applications, the integration server provides a flexible, manageable, and efficient way to provide tag and event information to any application that wants it.
- It implements business logic — determining when and where the event data is directed — giving more control over the data.
- It abstracts the Auto-ID architecture from back-end applications, which makes the tag data accessible to any back-end applications without the need for custom coding. In the prototype, the tag data is used both by the IMS and the order management system.

The prototype highlights how tag and event data from the Savant is just another data stream into the existing ERP application infrastructure.

Figure 8 illustrates the flow of events, and identifies the products and technology involved in the Sun iForce Auto-ID implementation. It demonstrates how RFID event data is combined with processes and data from back-end ERP applications to notifying a customer when specific items—cartons of razor blades—are shipped

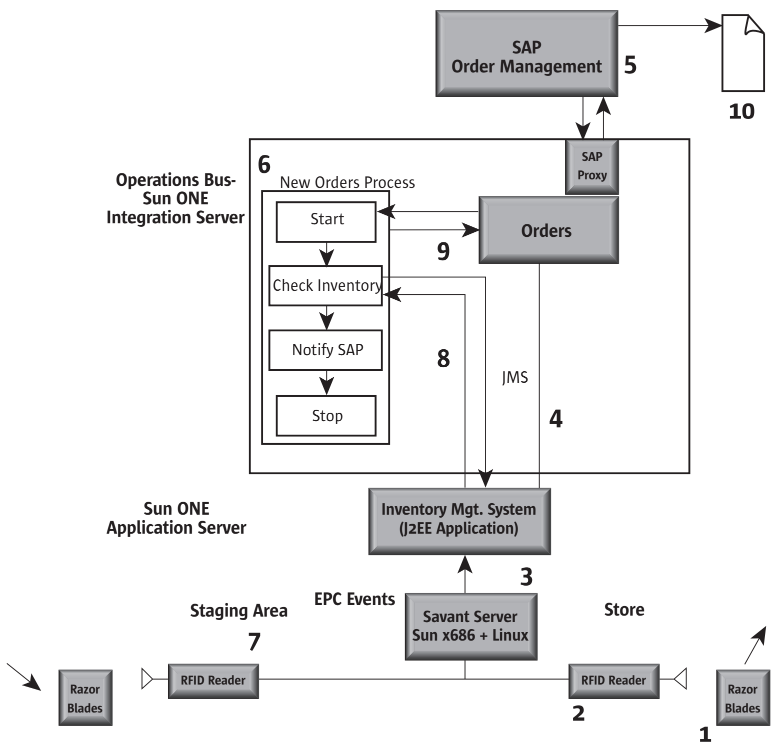

Figure 7: High Level Functional Architecture of ASN

1. Razor blades are removed from the shelf of a local retailer.
2. The RFID reader sends the signal to the Savant
3. Once a threshold is reached, an order for replacement stock is sent to the operations bus
4. The operations bus sends the order to the SAP application.
5. A new order for thirty cartons of razor blades is then generated automatically and can be seen on Pending Order screen of the SAP Order Management application. SAP sends this to the operations bus, in this case an instance of the Sun ONE Integration Server, Enterprise Edition.
   - SAP and the Integration Server are connected using a commercially available SAP connector.
6. The order triggers an integration server process: NewOrderProcess. This process instructs the IMS to place the order in the Pending Order queue. The order shows up in the Pending Orders screen of the IMS. The IMS checks for the thirty cartons in inventory.
   - The NewOrderProcess uses Java Message Service (JMS) to send the message from the integration server to the IMS.

- The prototype IMS application is a J2EE™ application implemented on Sun ONE Application Server 7.0

7. As the order is placed in the staging area, the IMS tracks the merchandise against the required amount. RFID tags on each pallet are read by RFID readers, and this information is forwarded to the Savant.
   - The Savant filters and smoothes the data, as required. The Savant also can be programmed to report on other thresholds, such as the last time a specific tag number was read.
8. When the 30 cartons are ready, the IMS sends a JMS message back to the NewOrderProcess within the integration server. The order is removed from the IMS Pending Order screen, and the order is now reflected in the IMS Completed Order screen.
9. The conditions for the new order have been met. NewOrderProcess communicates back through the SAP Proxy to the SAP Order Management application, indicating the order is complete, and will ship that day.
   - In the SAP Pending Order screen, the order is removed. The order appears in the SAP Completed Orders screen.
10. The SAP system sends an ASN message to the customer, telling them the order has been completed and will ship that day.

In the prototype, the Sun ONE Integration Server, Enterprise Edition is used to coordinate interaction between the ERP (SAP) and the IMS. Data translation as well as business process management is useful in optimally leveraging realtime supply chain information contained in the Savant with enterprise applications. Depending on the specific requirements either session beans or servlets can be developed for Sun ONE Application Server or Sun ONE Integration Server to address more complex integration scenarios.

Communication between system components is accomplished through technologies such as Java Messaging Service (JMS). The J2EE™ Connector Architecture (J2EE CA) can also be used to connect system components.

Future Directions

This proof-of-concept implements an end-to-end RFID-enabled Supply Chain Execution Infrastructure solution. It demonstrates the standards based integration layer architecture which automatically propagates information on product movement through the Enterprise Information System. This is running today at our Menlo Park iForce Center, and demonstrations can be easily arranged.

It is easy to envision additional applications for this powerful technology. Sun Microsystems is implementing solutions for other use cases. One example is a system that uses Auto-ID in the process of managing a product order between a manufacturer and retailer. As the order is picked and assembled by the manufacturer its accuracy is verified. As it is shipped, a dispatch notice is automatically sent to the retailer. As the order is received by the retailer, it is automatically verified. The future will see information automatically flowing along with (and ahead of) the physical movement of products and goods.

As an AUTO-ID Partner from the Beginning, Sun:

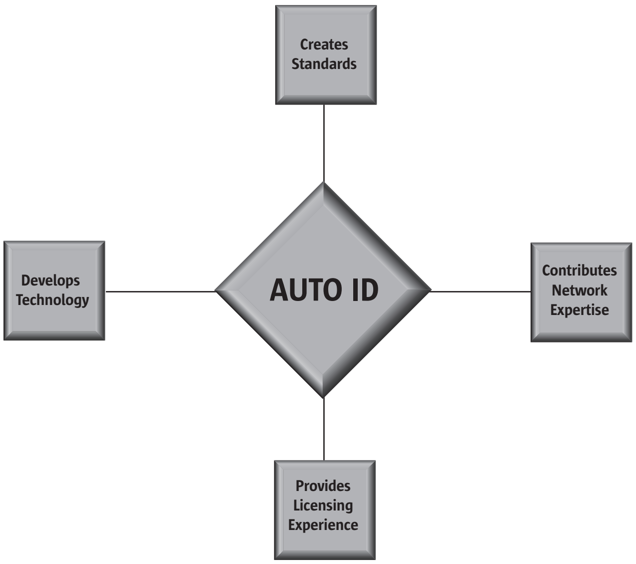

Figure 8: Sun is an Experienced Leader in Auto-ID Vision, Technology, Deployment

Sun Services

Sun Services experts have real-world experience installing the Sun ONE Content Delivery Server in complex operator environments, and can help deliver product upgrades or create customized integrated solutions. Together with partners, Sun Services provides a continuum of expertise, technology, and global coverage to assess customer needs, then implement and manage Sun Infrastructure Solutions to help organizations realize the full value of IT investments. Deep technology expertise, broad service offerings, and global experience serving enterprises make Sun Services the best choice for companies looking to reduce the time, cost and risk of transforming businesses with technology. Sun Services offers specific consulting, training and preemptive support. To help maximize service levels while minimizing costs, Sun Services gets solutions up and running quickly and efficiently, and can provide operational support and management capabilities.

Determining lead times for customized, wireless content delivery and integration is dependent on the technical discovery process. During this process, Sun Services meets with the customer to determine the scope of the project. Sun Services offers a well-defined, comprehensive engagement model.

Other resources include:

The SunTone℠ Initiative: Helping Achieve Service-Quality Goals

Sun Infrastructure Solutions give companies a solid solution design, architecture, and implementation. The final step in ensuring the quality of any solution is the operational management practices to meet service-level agreements. These disciplined practices, such as change management controls, are critical to service quality and availability. The SunTone™ Initiative certifies service delivery and management practices to validate quality against demanding service-level agreements. The program provides best practices for IT service management that helps companies streamline operational practices and drive efficiencies up front, while addressing people and process requirements to meet ongoing service delivery and management quality goals such as availability, reliability, response times, and security. What's more, SunTone professionals can help certify grid deployments, to help ensure they meet the highest levels of reliability and service delivery.

iForce℠ Solution Centers

Sun's iForce℠ Solution Centers enable customers to simulate real-world production environments in a risk-free setting. At these centers, customers work jointly with experts from Sun and appropriate iForce partners to build and test proof-of-concept demonstrations and run pilot programs. The iForce Solution Centers help organizations identify areas for improvement, validate their architectures, and accelerate the deployment of high-quality Infrastructure Solutions.

iForce Partners

Sun Infrastructure Solutions can be extended through the addition of components, applications, and services from Sun's market-leading, best-of-breed iForce partners—ISVs, OEMs, service providers, integrators, system providers, solution resellers, and VARs of Sun systems and platform-based solutions. Sun understands the value of partners and works with them to create, market, and deliver network computing solutions that help reduce costs and increase customer satisfaction.

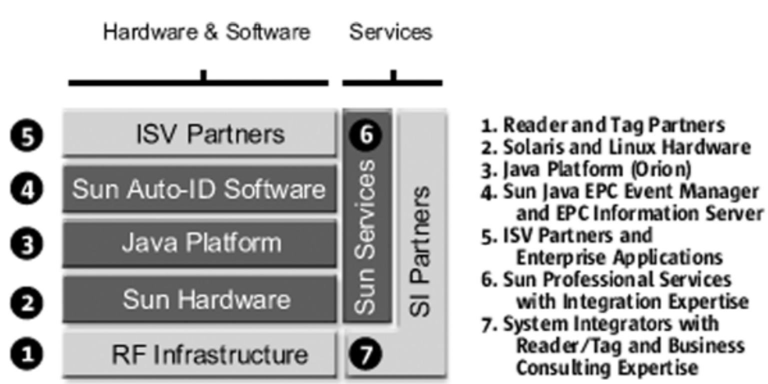

Figure 9: Sun's Recommended Software Solution Components

Conclusion

Auto-ID technology holds the promise to create significant cost savings in supply chain management, and the potential to create revolutionary applications and services. The adoption of standards will allow companies in a supply chain to send information along with products, and monitor this information along the way. The Auto-ID standards, with the support of nearly 100 sponsoring companies, are creating a cost-effective way to make the supply chain more efficient.

Sun has been an early sponsor of the Auto-ID Center, and provides support with technology committee leadership. Furthermore, Sun has participated in early trials with leading consumer goods manufacturers and retailers.

While creating a cost-effective way to generate and collect product data is a significant challenge, so is integrating it into the EIS and ERP systems and repositories. Modern businesses rely on many back-end applications to run their business, and many of these will need access to the data feeds from Savants distributed across company operations.

One of the findings from Sun's efforts is the utility of an operations bus in the deployment. Advantages include making the data generally available to virtually any enterprise application, overcoming the burden of point-to-point connections, and the ability to create new services from existing applications—such as the iForce ASN prototype. When an integration server such as the Sun ONE Integration Server, Enterprise Edition is used as an operations bus, information can be transmitted in a secure and standardized way from every Savant. In addition, business logic at this layer typically can be created or modified by business analysts, without the need for low-level programming, enabling faster reaction to business opportunities.

The iForce ASN reference implementation is highly complementary to the Auto-ID Center's standards and specifications—it's built on open standards with open APIs, uses Java technology, combines the efforts of best-of-breed partners to create value-added solutions.

Sun's vision—that even basic items will be connected by "The Internet of Things"—is coming closer to realization. The standards developed by the Auto-ID center and its members are enabling this dramatic change and the productivity leaps that go along with it. Sun's Enterprise Auto-ID Architecture and its implementation in the iForce project described above is proof of their leadership in this market. Sun and its iForce partners will help your enterprise leverage the investment that has been made in existing systems while enabling you to take advantage of the rapid emergence of Electronic Product Code standards being developed by the Auto-ID initiative.

> "Every single object manufactured will be created with a wireless address and will carry an identifier"
> –Bill Joy, Sun Microsystems—1999

Chapter 6

More Information

To learn more about the Sun's Auto-ID technology and products, see the Web sites below. Organizations can also contact Sun or its partners to learn how we can help build competitive advantage with solutions that match business goals and processes to technology solutions.

sun.com/autoid	Sun Auto-ID
sun.com/erpscnetwork	iForce ERP/Supply Chain Network
sun.com/iForce	Sun's iForce Initiative
java.sun.com	Java Technology and Standards
sun.com/sunone	Sun Open Net Environment (Sun ONE)
www.autoidcenter.org	Auto-ID Center
www.projectliberty.org	Project Liberty Alliance

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 800 786-7638 or +1 512 434-1577 **Web** sun.com

We make the net work.

Sun Worldwide Sales Offices: Argentina: +5411-4317-5600, Australia: +61-2-9844-5000, Austria: +43-1-60563-0, Belgium: +32-2-704-8000, Brazil: +55-11-5187-2100, Canada: +905-477-6745, Chile: +56-2-3724500, Colombia: +571-629-2323 Commonwealth of Independent States: +7-502-935-8411, Czech Republic: +420-2-3300-9311, Denmark: +45 4556 5000, Egypt +202-570-9442, Estonia: +372-6-308-900, Finland: +358-9-525-561, France: +33-134-03-00-00, Germany: +49-89-46008-0 Greece: +30-1-618-8111, Hungary: +36-1-489-8900, Iceland: +354-563-3010, India–Bangalore: +91-80-2298989/2295454; New Delhi: +91-11-6106000; Mumbai : +91-22-697-8111, Ireland: +353-1-8055-666, Israel: +972-9-9710500 Italy: +39-02-641511, Japan: +81-3-5717-5000, Kazakhstan: +7-3272-466774, Korea: +822-2193-5114, Latvia: +371-750-3700, Lithuania: +370-729-8468, Luxembourg: +352-49 11 33 1, Malaysia: +603-21161888, Mexico: +52-5-258-6100 The Netherlands: +00-31-33-45-15-000, New Zealand–Auckland: +64-9-976-6800; Wellington: +64-4-462-0780, Norway: +47 23 36 96 00, People's Republic of China–Beijing: +86-10-6803-5588; Chengdu: +86-28-619-9333 Guangzhou: +86-20-8755-5900; Shanghai: +86-21-6466-1228; Hong Kong: +852-2202-6688, Poland: +48-22-8747800, Portugal: +351-21-4134000, Russia: +7-502-935-8411, Saudi Arabia: +9661 273 4567, Singapore: +65-6438-1888 Slovak Republic: +421-2-4342-94-85, South Africa: +27 11 256-6300, Spain: +34-91-596-9900, Sweden: +46-8-631-10-00, Switzerland–German: 41-1-908-90-00; French: 41-22-999-0444, Taiwan: +886-2-8732-9933, Thailand: +662-344-6888 Turkey: +90-212-335-22-00, United Arab Emirates: +9714-3366333, United Kingdom: +44-1-276-20444, United States: +1-800-555-9SUN OR +1-650-960-1300, Venezuela: +58-2-905-3800, Or Online at sun.com/store

SUN™ THE NETWORK IS THE COMPUTER

LFC X.X Printed in USA 09/03 XXxxxx/XK