# The Knowledgeable Utility and SISCO's Utility Integration Bus

Do you have a vision for a utility which can make business decisions based on all the information available to it?

Do you have a vision for a utility that is responsive to the needs of customers, shareholders and employees?

Do you have a vision for a utility in which E-Commerce is fully and securely integrated into its infrastructure?

This vision is the essence of the competitive utility. To achieve this vision your enterprise must be able to make decisions based on a rapidly fluctuating energy market, information about the customer, and internal operational data. To achieve this vision large amounts of data must be integrated and rapidly analyzed. To achieve this vision you must begin to look at Information Technology (IT) as more than just a back office operation. Achieving this vision will build the *Knowledgeable Utility*.

The knowledgeable enterprise is a place where managers can create views of and analyze all available information regardless of where the data is stored. A place where IT managers are full partners in the management decision making process. Only by understanding the current energy market and ones own ability to supply or transmit power can one profitably take part in the deregulated environment. This white paper describes "The big picture" and how SISCO's Utility Integration Bus (UIB) can help you create a knowledgeable utility.

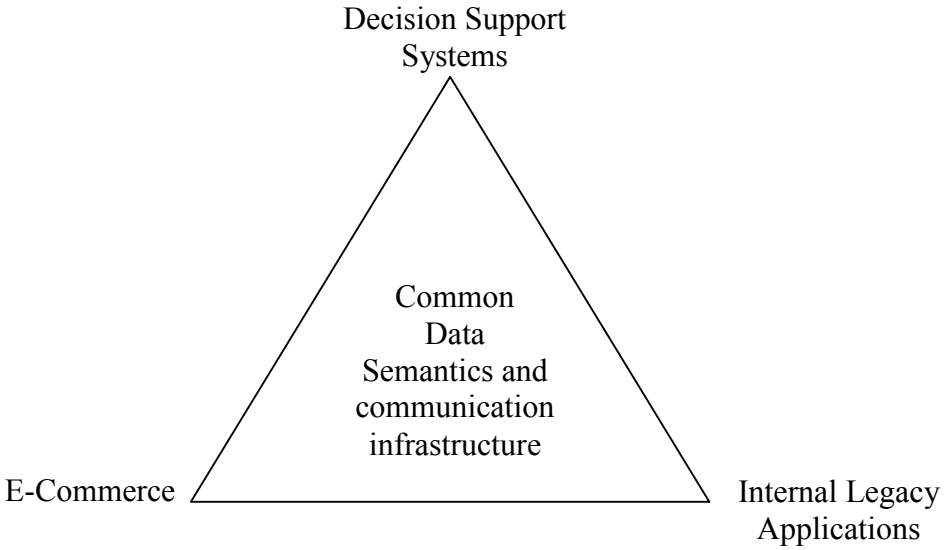


Knowledge is data that has been analyzed and put into a meaningful context. In order to create a knowledgeable utility financial, managerial, and operational systems must be fully integrated. Decision Support Systems (DSS) help create knowledge by analyzing customers/suppliers related data gleaned from E-Commerce applications as well as operational data from enterprise applications. A communication infrastructure facilitates the flow of data to analysis tools. A useful E-Commerce front end must exchange information with back end systems and is tuned to market conditions based on knowledge from Decision Support Systems. Integration of legacy applications requires a communication and process infrastructure. Consequently, one can not create a knowledgeable enterprise with only Decision Support, Integration and E-Commerce tools alone. This is a case where the total can be greater that the sum of the parts. To harness the maximum power of these subsystems a utility must have an infrastructure to support standard communication methods as well as the management of the meaning (semantics) of the data used though the enterprise. The communication infrastructure must be able to integrate all part of the enterprise. Data semantics must be able to be organized according to managerial and operational needs. Only in combination can a communication and semantic infrastructure maximally facilitate integration and the exchange of shared knowledge. Only in this way can these Decision Support, E-Commerce, and legacy applications be maximally leveraged.

## The Utility Integration Problem

Today's utility IT environment is truly heterogeneous. Some of the more significant features of this mix include:

- Many computing hardware platforms
- Many operating systems
- Mainframe/client-server/web-based systems
- Many component technologies (CORBA, DCOM, Enterprise Java Beans)

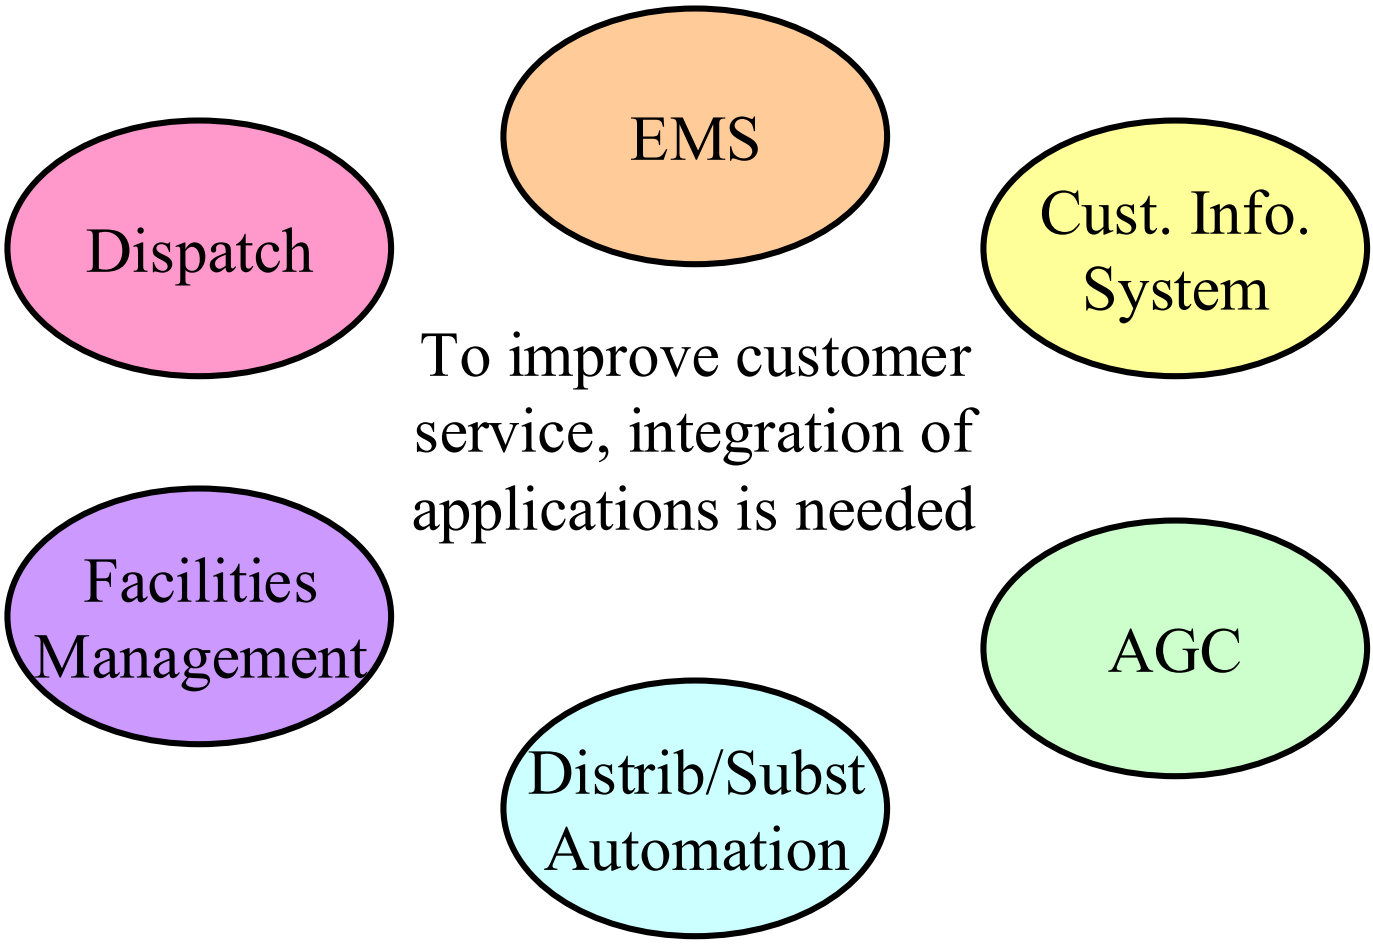


Furthermore, the utility of today has several sources of "real-time" process oriented information that allows the control and monitoring of key devices, measurements, and subsystems. These include: Supervisory Control and Data Acquisition (SCADA) systems, Energy Management Systems (EMS), Distribution Facilities Management Systems (DFMS), Automatic Meter Reading (AMR) systems, and other sources of "real-time" information. Several of these systems may need to exchange information with business applications and desktops or with other process oriented systems. For example within the control room, an EMS obtains power system information from the SCADA system and then commands the SCADA system to perform control on its behalf. Availability of process information outside the control room would allow further integration of operational data into business applications such as Work Management Systems (WMS). A WMS contains information about workorders or asbuilt designs that would be valuable to dispatchers in the control room. WMS and other systems like it

typically employ databases that have been configured with customer and SCADA system information— any or all of this information may be valuable to another application.

In the past, utilities have undertaken short terms solutions which has resulted in many point to point links as shown in the diagram below.

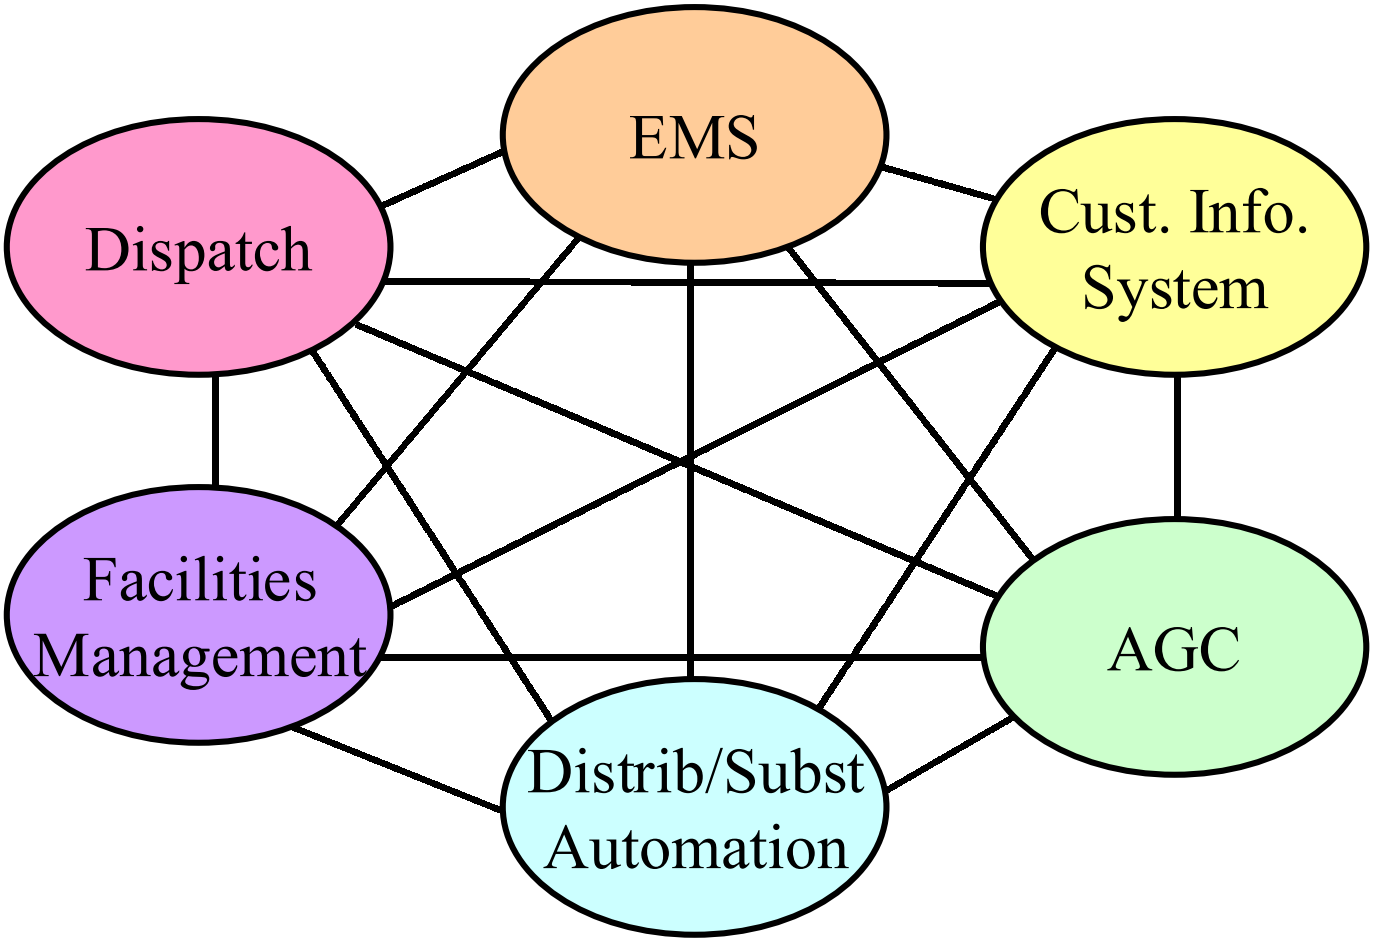


Integrating a utility's various systems requires a method that is technology neutral and that does not require existing applications to be disturbed. Recently a new type of middleware has emerged whose sole function is to link legacy applications via the exchange of messages. This type of software requires building wrappers for existing applications and can provide a standardized way to perform integration. In general, they are responsible for exposing the functionality of the wrapped application in a way that is compatible with a common methodology. They also perform translation between object models. The most significant benefits of this approach are that it is technology neutral and that legacy applications do not need to be rewritten.

One problem with using type of approach is that adapters can become large. A significant amount of code is required to abstract away different calling conventions, design assumptions, data models, and mismatches in assumptions about business processes. As the size of adapters increases, so does the cost of writing and maintaining them.

Another problem with this type of software is that it only provides tools for inter-applications communication and does not facilitate the creation and management of common data semantics. Thus while integrating existing applications using these tools solves operational problems, in fact the real value of doing integration can be that it facilitates the creation of higher level data analysis applications. One needs to not only create a communication infrastructure to automate the exchange of data, but also to establish common data semantics. In this way data in existing systems can be can become shared knowledge.

## Application Integration with the UIB

Fundamentally, application integration with the UIB involves looking at the big picture. However, integration may encompass data from a large or small set of applications. One does not need to undertake a major project that requires many months to complete. The issue here is the development of a long term enterprise wide integration strategy so that a small integration project does not become just another slightly larger island of automation. Thinking at the enterprise level while integrating at the department level minimizes risk and maximizes the chances for long term success. Part of this enterprise view is the understanding of enterprise data semantics and the business decision making process.

The understanding of data semantics requires a unified data model. The coalescing of an enterprise's many data models into a smaller more rational set whose purpose is to enable decision making is often called data integration. Data integration is some what different from most programming tasks in that the goal is not necessarily to add new features, but rather to link and expose existing data while minimizing reprogramming. The creation of an integration bus is inextricably linked to the creation of shared data models.

In order for a data model to be used by multiple applications, its semantics must be understood and managed. The commonly accepted way to manage data semantics is by describing what, where, and how data is used in a metadata repository. Metadata is data about data. A metadata repository serves as a central point of control for data semantics, providing a single place of record about information assets across the enterprise. It documents where the data is located, who created and maintains the data, what application processes it drives, what relationship it has with other data, and how it should be translated and transformed. This provides users with the ability to utilize data that was previously inaccessible or incomprehensible. A metadata repository also allows all applications to be integrated to communicate using a unified data model instead of creating multiple point to point links. Lastly, a central location for the control of metadata

ensures consistency and accuracy of information, providing users with repeatable, reliable results and organizations with a competitive advantage.

To support both DSS, E-Commerce, and operational processing, the UIB manages a wide variety of metadata, including:

- Business definitions
- Business formulas
- Business relationships
- Table and column descriptions
- Attribute definitions
- Data transformation descriptions
- Data routing information
- History of changes

The DSS uses this metadata to put so users can drill down and discover the context of analysis results. Without knowing the context of an analysis one can never be sure what any result means. The communication infrastructure uses this metadata to determine how to exchange data. Without knowing how data is routed and exchanged, application integration becomes impossible.

For example, the UIB allows utilities to fully leverage common data models such as the Electric Power Research Institute's (EPRI) Common Information Model (CIM). The CIM is a data model that represents all the major objects in an electric utility enterprise typically contained in an Energy Management System. The CIM is being accepted and implemented by a growing number of organizations as an ideal model for building new utility operations applications and integrating legacy applications. Applications that have been integrated or are planned for integration with the CIM include:

- SCADA
- AGC and Economic Dispatch
- Interchange Scheduling
- State Estimation and Powerflow
- Operator Training Simulator
- Asset Management

The CIM is primarily organized to meet the needs of operational systems. While an operational perspective is fine for routine needs, the management of the Knowledgeable Utility really needs an infrastructure for creating summaries and performing analysis in order to identify trends, challenges, and new opportunities. The UIB provides multiple views of the data model to support both management and operational needs. Thus, the UIB provides a semantic infrastructure so that E-Commerce and Decision Support applications can be easily built upon existing legacy systems.

Besides providing a way to manage common data models such as the CIM, the UIB provides a common way to automate communication. The UIB acommplishes this by automating the exchange of eXtensible Markup Language (XML) messages. Just as HyperText Markup Language (HTML) has become the universal language of the Web,

businesses have sought a similar language for describing business data. XML has been adopted by the World-Wide Web Consortium (W3C) and is rapidly becoming the preferred format for exchanging complex businesses data internally as well as between E-Commerce applications. Similar to HTML, XML allows the designer to create custom tags and describe how they are used and thus provides the facilities to create self describing messages. This capability is independent of how transport mechanisms, calling conventions (the order in which parameters are passed as well as how data is returned), and data formats. This significantly reduces the size and complexity of legacy application wrappers. XML- formatted business data offers standard and extensible information formats or packages with which to exchange information internally and with other businesses. The UIB uses XML messages so that E-Commerce can be more easily integrated with back end systems as well as to take advantage of XML standardization and extensibility.

But companies still need a reliable mechanism to send and receive XML packages. One could use CORBA or DCOM to link legacy applications, but these technologies require a common security domain context, function calling convention, binary data types, and way of locating and activating remote applications. Additionally, CORBA and DCOM require that server applications must be ready to service a request when the client wishes. Thus CORBA and DCOM are better suited to assembly of tightly coupled components. To use a post office analogy, no one waits at the front door for the postman to arrive before mailing a package. Mailboxes provide a convenient method for storing letters until a mail truck comes along to pick up the mail and deposit the received mail. One could use email, but email has not been designed for efficient automation. As mentioned above, existing message oriented middleware products help link applications. In general, these software products include a message broker. With message broker technology, a business application can send business messages to a broker message queue for later delivery. The messages are then picked up by the message broker and dispatched to other internal or external applications. Message brokers facilitate location and technology independence and have proven to be the best way to link loosely coupled legacy applications

In addition to message brokering, these middleware products often include the ability to automate business processes (data transformation and workflow). In doing so, they provide a single point of control for managing information flow across multiple applications. For example, the generation of a bill can be automated by creating a script that first collects meter and customer data, then sends the information to a billing application and lastly routes the bill to an application for presentment to the customer. In between these steps, message data may be manipulated so that it matches the internal data model of these applications. The UIB helps integrate business process metadata to create a centralized maintainable enterprise view of data processing.

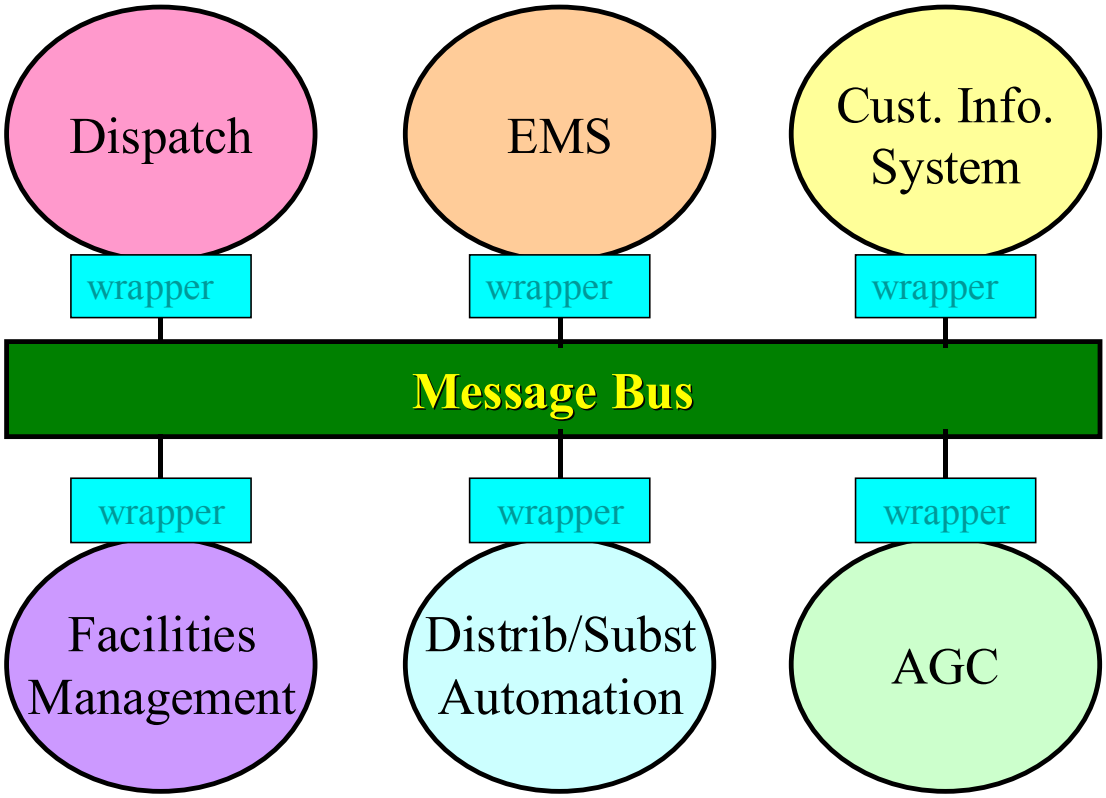


Using XML messages routed via a workflow enable message broker vastly simplifies application wrapper complexity. While wrappers still need to perform potentially complex translation between object models, this too can be greatly simplified. Ideally, applications would export their metadata into the metadata repository using a standardized mechanism so that the actual transformation code could be automatically generated using a visual tool. In fact, several standards groups are working on defining common ways to export metadata via XML. Companies such as SAP, Baan, Peoplesoft, and JD Edwards have committed to adding standardized metadata export functionality. While at this point the UIB does not include this tool, even a manual process is greatly eased by a standard metadata export file format. Lastly, as the CIM becomes more widespread as the standard object model for control center application integration, it is likely that utility specific vendors will provide standard mapping between their proprietary object model and the CIM.

The infrastructure described above system can extensively use a publish/subscribe message passing paradigm. In this case, data consumers may not know who is the data providers are and visa versa. Consequently, there must be some component that knows if all events are being processed. The UIB maintains the status of all UIB application

components. Additionally, since the UIB is aware of where data originates, the UIB can monitor data availability. Thus the combination of a metadata repository and messaging provides a powerful way of performing real time monitoring and analysis of information availability and flow.

## UIB Architecture

The diagram below illustrates a business application or automation agent with two different sets of interfaces available to it. The first set of interface definitions represent

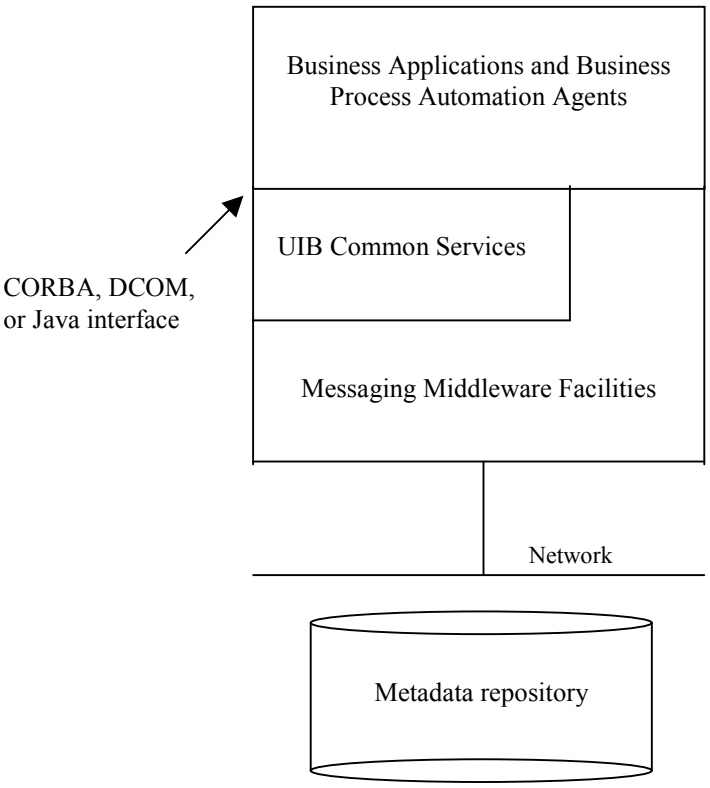


**General Architecture**

how a business application is to invoke one or more of the UIB's common services. The second set of interfaces is provided by a middleware provider. The middleware is comprised of two parts: A messaging service provider such as IBM's MQ Series and a business process (data transformation and workflow) engine. Although the actual interfaces depend on the component technology used (CORBA, DCOM, Enterprise Java Beans), the funtionality of UIB services are independent of component technologies and underlying messaging middleware. Since the UIB encapsulates the underlying middleware features business applications are free to use either interface. Business applications access UIB services as a CORBA, DCOM, or Java component.

## Conclusion

In the past utilities were highly regulated companies producing a stable stream of earnings, but not rapid growth. The reorganization of utility business entities makes this scenario much less certain. Furthermore, the combination of more dynamic energy markets, and the growing importance of E-Commerce is driving the greater use of Decision Support Systems and application integration tools. Key to successful implementation is an enterprise level architecture describing how knowledge is shared. For some utilities, this quest has lead to a reexamination of how of how IT fits into management. As IT becomes more integrated into the enterprise, IT becomes a full partner in the management decision process. Activity to supply these enterprise level capabilities is starting to be aggressively pursued by all major software vendors. It is critical that IT managers to understand the issues at hand as they adopt enterprise strategies and that they be in a position to evaluate what set of vendor products are appropriate to their situation. The UIB has been created to maximally leverage existing standards and products. Today the UIB is the only product to fully take advantage of all of these technologies and provide the crucial missing pieces.