Interoperability Platform

A Bentley White Paper

A.B. Cleveland, Jr.
Senior Vice President
Bentley's Applied Research Group

Bentley Sustaining Infrastructure

www.bentley.com

Executive Summary

Sustaining our society and our environment requires infrastructure. While infrastructure often becomes permanent, every infrastructure project is temporary. Driven by new technologies, new requirements, new partners, and focused on infrastructure assets bound to a specific geographic location with its own demands, each project is unique. But all projects face a number of common challenges, including increasing demands for lower cost and shorter schedules; project teams that consist of multiple, globally distributed organizations; and 21st-century demands that relate to security and sustainability. Fundamental to the success of any infrastructure project is information, and how that information is created, shared, distributed, delivered, and applied. Yet, design teams must execute each infrastructure project with an information-technology environment that includes a host of software applications from several vendors, each with differing information models and file formats, and is typically interspersed with a significant amount of manual effort. This is particularly problematic as information must flow across functional, organizational, and lifecycle phase boundaries.

The lack of an interoperable platform among its information technology systems often frustrates project teams and impedes their ability to execute projects more successfully. Moreover, project teams are equally frustrated by typical approaches to address the lack of interoperability – monolithic, centralized systems incompatible with the dynamics and inherent uniqueness of each project, and point-to-point information sharing between applications, which seldom goes beyond importing and exporting file formats and does not scale. Neither of these two approaches provides the flexibility demanded by the dynamics of infrastructure projects, nor allows project teams to opportunistically leverage project information to their strategic advantage. As a result, the AEC community requires a new approach that enables infrastructure projects to achieve interoperability – one that allows project teams to execute more effectively through information technology; acknowledges and embraces the dynamics and uniqueness of every infrastructure project; commits to supporting current and future infrastructure industry standards; and provides the opportunity to further leverage project information and increase the return on information technology investments.

Bentley Systems offers such an approach through its *interoperability platform*. Bentley developed this platform in response to the unique requirements of infrastructure projects. Based on a federation strategy, it enables infrastructure projects to achieve interoperations among its information technology systems by creating an interoperable project environment. In order to achieve an interoperable project environment, Bentley's interoperability platform includes the following key characteristics:

- *Embraces existing applications.* The interoperability platform offers an increasing array of off-the-shelf connectors, which enable heterogeneous information sources to connect to the interoperable project environment. Software toolkits make it possible to create new connectors for additional information sources. Each source requires only a single connector, regardless of the number of other applications connected to the interoperable project environment.

- *Synchronizes project information.* Information can be synchronized across multiple applications, ensuring that all applications appropriately reflect the current state of shared information. This information can be aggregated, transformed, and distributed to other applications and systems based on project rules. A key element of this capability is the reliance on dynamic schema and the mappings between schemas, which are maintained externally to the application software, where they can be dynamically adjusted to adapt to changing project conditions and configurations.
- *Generates integrated views of project information.* Through the same flexible capabilities and underlying technology used for synchronizing project information, the interoperability platform can aggregate and transform project information into integrated "views" of the current project state that span multiple applications and information sources. Project teams achieve this without resorting to inflexible or impractical solutions, such as a centralized database. These views provide increased visibility into the project and a strategic value that is impractical or impossible to obtain with other approaches.
- *Creates dynamic deliverables*. Based on the integrated views into the project, the interoperability platform can create dynamic deliverables in a variety of forms and formats. These dynamic deliverables are provided in multiple formats ranging from integrated intelligent models and drawings, industry standards such as ISO 15926, intelligent web streams, intelligent publishing formats like PDF, and intelligent paper deliverables.
- *Provides powerful client applications for extracting value from project information*. The interoperability platform provides a robust set of client applications (thick clients, web clients, information portals) to consume the dynamic deliverables. It also provides a robust set of tools to create new applications that extend beyond visualizing and commenting on project information. This expanding array of applications can be used for design review, analyzing the state/status of project execution, integrated project deliverables, construction, operations work orders, and multiple types of analysis, including safety, performance, and security, among others. These applications "consuming" the dynamic deliverables can include those that provide feedback (e.g., redlines) to the content creators, as well as those that use the dynamic deliverables as the starting point for the creation of new content.

An interoperability platform such as the one Bentley provides can deliver broad, strategic benefits to infrastructure projects, including greater project-level productivity; increased flexibility and resilience for infrastructure project enterprises; increased reuse of project information; increased visibility into project information content; and creation of new knowledge that can be reused on subsequent projects. Further, an interoperable environment enables infrastructure projects as well as its individual participants to be better positioned to identify, create, and leverage new opportunities.

This paper explores the issues surrounding interoperability, particularly as it relates to achieving interoperations across functional, organizational, and lifecycle phase boundaries. It describes different approaches to achieve interoperability and provides an outline of Bentley's vision for an interoperability platform that drives Bentley's technical approach to implementing an interoperability platform into the hands of our users.

The Case for Interoperability

> "The average case...does not exist at all, because every case, as soon as it really occurs, at once becomes a quite special case, and sometimes it is absolutely unlike anything that has ever happened before."
>
> –Fyodor Dostoevsky, 1864

Projects are the engines that drive an infrastructure asset's lifecycle. Through projects, infrastructure assets are created, expanded, improved, adapted to new technologies, modified for new purposes, and decommissioned. Executing infrastructure projects – regardless of the lifecycle phase – has always brought with it significant challenges. The project enterprise is always temporary, brought together quickly and for a short time to build, modify, or decommission the infrastructure asset, and then disbanded. In fact, it is extremely unlikely that the exact configuration of a project enterprise will be duplicated in a later project. Advancement in technology, new suppliers, and new requirements can significantly impact the project deliverables in terms of new or modified infrastructure, which inevitably changes the configuration of the project enterprise. This dynamic is reinforced because every infrastructure project is bound to a specific location on the earth. As such, it is subject to local regulations, the local political and economic environment, local practices and methods, local suppliers, and a local workforce. These factors make every project unique.

Recent trends have added new pressures that further increase the complexity and difficulties of project execution. Some of these include an increasing skilled-labor shortage; emerging sources of expertise, technology, and manufacturing capability in the developing world; and growing infrastructure requirements in the developing world, which have led to an increase in globally distributed projects. In certain domains, "mega" projects have become the norm. Other pressures include shortened schedules and faster time-to-market demands. Add to these the 21st-century demand for greater security and sustainability.

To meet these challenges more effectively will require a response from all dimensions of the infrastructure industries. This paper addresses the role of information technology systems, specifically the issue of interoperability during project execution, as well as the continued use of the content created by infrastructure projects.

The International Standards Organization (ISO) defines interoperability as the:

> *Capability to communicate, execute programs, or transfer data among various functional units in a manner that requires the user to have little or no knowledge of the unique characteristics of those units.*

To put it in layman's terms – "All of the *software* used on the project should *work together*, and making it work together should be *easy*." Of course, the objective is not just to implement systems that are interoperable – the objective is to realize actual interoperations among our information technology systems. This is simple to say, but quite difficult to realize in practice. As pointed out in a 2004 report by the National Institute of Standards and Technology (NIST), a lack of interoperability costs the infrastructure community $15.8 billion annually in the United States alone, which amounts to between one and two percent of the industry's revenue. This estimate includes only costs and does not begin to address the impact of the *lost opportunities* due to a lack of interoperability. This paper will explore issues surrounding interoperability, describe some different approaches intended to achieve interoperability, and provide an outline of Bentley's vision driving our implementation of an interoperability platform.

Information Technology and Interoperability

If projects are the engines that drive an infrastructure asset's lifecycle, then information is the fuel. Projects succeed or fail on how effectively information is created, shared, applied, and distributed. If you accept this premise, then the obvious conclusion to draw is that information technology plays a central role in successfully executing infrastructure projects and meeting the challenges outlined in the previous section. In spite of the significant advantages we already realize from information technology, the key to unlocking the full value of information technology is to solve the interoperability problem.

> "Projects succeed or fail on how effectively information is created, shared, applied, and distributed."

In terms of today's interoperability problems, do these scenarios sound familiar?

- Multiple application software systems, each storing its data in a different format with different mechanisms for accessing that data, are used on a project;
- Each application has an independent development path, evolving in response to its specific user community, and providing functionality in a narrowly focused domain;
- The application systems are a mixed bag of systems – some developed in-house, some purchased commercially, and some falling in between;
- Sharing data between the systems, particularly across discipline or department boundaries, is usually achieved at best – if it doesn't simply default to an error-prone manual transcription process – with specialized point-to-point interface software;
- There is an integration scenario in place, but each time an application is added, replaced, or upgraded, some level of software development is required to accommodate the change in the application mix, often at a cost that cannot be borne by a single project;
- Engaging in a project with a new partner, under a new business model or with a new contractual arrangement requires a major change, if not replacement, of your internal systems and the associated interoperability approach;
- You work downstream from design in the infrastructure asset's lifecycle (e.g., construction, operations, etc.) and the electronic data available to you from the upstream organizations is not in the format you need, is not the right scope for your use, and does not integrate easily with the information unique to your process.

This is certainly not an exhaustive list of the "pain points" related to a lack of interoperability, but it does represent the kind of issues each infrastructure project faces. These examples are much closer to the *rule* than the *exception*. As problematic as the long list of pain points caused by a lack of interoperability might be to an individual project, the impact can be far more strategic. It also limits opportunity for much-needed innovation in the form of new work processes, new business models, new contractual models, new delivery mechanisms, and so on, simply because it is too difficult or costly to opportunistically adapt our information technology systems to new configurations. How many times have new business configurations not been pursued, or worse, failed when attempted, because the information technology system wouldn't support them, or was not flexible enough to respond in time to support a project? These lost opportunities reflect the real cost of a lack of interoperability.

So we must ask: What makes interoperability so difficult to achieve? If it were simply a matter of wiring applications together using low-level standards – XML, HTML, SQL, etc. – the interoperability problem would have been solved long ago. But *connectivity* is not the same as *interoperability*. Interoperability requires that we not only connect disparate applications but that one application appropriately comprehends the subset of semantically rich information created by another application relevant to its domain. The differences in these representations might be the biggest roadblock to achieving interoperability.

> "But connectivity is not the same as interoperability. Interoperability requires that we not only connect disparate applications but that one application appropriately comprehends the subset of semantically rich information created by another application relevant to its domain."

Sources of Difference

Between any two applications there are multiple dimensions to how information is represented, each with its own implication. Many of these differences result from how people use information as well as the software application systems that actually create and manipulate that information. These sources of difference fall into the following categories: data formats, semantics, relationships, scope, and level of detail.

Data formats

Information that, at the project level, is considered to be common between applications can differ in format in a number of ways. The most obvious difference is in the method of actually storing the data — SQL vs. Oracle, binary vs. ASCII, image vs. database, DGN vs. DWG. In fact, even the "human-readable" format can vary. In one form, the system name precedes a part number with a project number, while in another form it does not. In one form the system name includes dashes, while in another form it does not. In one form the size is expressed in feet, while in another it is expressed in meters. Those are just a few examples.

Semantics

No two programs or databases seem to refer to the same thing with the same name. Thus, when dealing with different systems, we may constantly be mapping between one name and another. Technically, this is a very straightforward issue to solve. However, if this were the only problem, differences in electronic representation of information would be nothing more than inconvenient.

Relationships

Relational database technology and the mechanisms for relating data within this context are well understood. However, the relationships between information elements, as well as the components and documents they represent, can be complex and indirect. In order for relational technology to work well, the information might need to be restructured, often with considerable difficulty and compromise. It is often impossible to construct geometric relationships, such as connectivity, with simple property value manipulation. However, in order to take full advantage of the electronic information, complex relationships and dependencies, both explicit and implicit, must be represented and accessible.

Scope

The scope of a collection of project information is particularly important as the information passes across any of the significant boundaries – engineering to construction, construction to operations, etc. – that exist within all projects. Within the organization that creates the information, such as engineering, there is a natural scoping of the information to be fairly consistent with the scope of the activities being performed. Coupling information between separate domains is much looser than coupling information within a single application or integrated workgroup. When that information is sent downstream, however, the recipient's requirements, with respect to how the information is grouped and segmented, might be quite different than those of the originator.

> "Corporate knowledge today still exists largely implicitly and in multiple forms. Often, corporate knowledge is not even recognized as such."

Level of Detail

Differences in the level of detail that various information users require can cause great difficulty when sharing information. Let's take, for example, a piece of mechanical equipment. Everyone who deals with the existence of that equipment – the mechanical engineer who defines its existence, the piping designer who must connect pipe to it, the purchasing agent who must expedite it, the electrician who must wire it, the maintenance engineer who must service it – requires varying levels of detail in the information they need about the equipment in order to do their job.

Information Technology and Knowledge

In spite of the significant use of information technology in infrastructure projects and the vast amount of information available in electronic form, most organizations do not *electronically* capture, manage, and apply their "corporate knowledge" to infrastructure projects. Corporate knowledge today still exists largely *implicitly* and in *multiple forms*. Often, corporate knowledge is not even recognized as such. The whole issue of managing corporate knowledge episodically receives attention within the largest organizations, but there is not as yet any overarching paradigm or methodology that has been pervasively adopted. The context of this paper focuses on semantic knowledge. In simple terms, *semantic knowledge* can be viewed as the business rules for executing a project. It establishes how information is managed, transformed, and distributed through the project from inception to completion. Today, this type of semantic knowledge typically exists in several forms.

Semantic knowledge is often represented in the form of procedures, guidelines, and standards that define the rules for executing a project, conducting business with suppliers and partners, and so on. This type of knowledge is often distributed in printed or graphic form such as paper, word processing documents, PDF files, and drawings. In a global operation, *consistent* and *timely* application of the knowledge in these forms is difficult. Even more problematic is the case of knowledge that comes from outside the company, such as supplier knowledge, since it is further removed from the internal mechanisms, such as they exist, for representing knowledge. Even in cases where central repositories for reference information have been established, there is almost exclusive reliance on the *individual* project-team member to access the knowledge and apply it appropriately and consistently with the rest of the organization.

Perhaps the most prevalent and effective location where semantic knowledge is stored today, particularly from a broad interoperability point of view, is "in the heads" of key project personnel. These individuals are often required to make critical decisions based on the information distributed throughout a multitude of application systems and databases. In order to make a decision, they must first retrieve the relevant information and then mentally integrate, order, and interpret it. But, relying solely on the implicit (and perhaps contradictory) knowledge within the minds of people can lead to decisions being made with information that is incomplete, out of date, or out of sync. This knowledge, while extremely valuable, is repeatable and reusable *only* if the same people work on the next project.

> "An organization's corporate knowledge is the resource that differentiates it from its competitors and forms the basis of the value it delivers to its customers."

Typically, the only place semantic knowledge is explicitly represented in an *actionable* form is within the multitude of software applications used in the project. This knowledge can, in fact, be easily reused simply by using the same applications on the next project. However, knowledge that is encoded in the form of compiled software suffers from a number of limitations:

- The knowledge becomes essentially "frozen" within the software. Extending or modifying the knowledge requires access to the application source code, which is certainly not readily available, if at all, for commercial application systems.
- Distributing the knowledge requires distributing the entire application since the knowledge component of the application does not exist independent of the application itself.
- The knowledge loses all "immediacy" in that the *control* of the knowledge does not reside with the end user, but with the owner of the source code.
- Lastly, most project applications are aimed at specific functions within a certain domain and do not address the broader issues of information flow across all the project's organizations and phases. They do not have a broad interoperability point of view.

An organization's corporate knowledge is the resource that differentiates it from its competitors and forms the basis for delivering value to its customers. Today, managing, applying, and distributing that knowledge is largely a manual process in spite of operating within a largely electronic environment. An organization achieves a significant advantage when it can explicitly represent its key knowledge and then distribute and apply that knowledge instantly and consistently to information that might be distributed around the world in a multitude of electronic forms.

The Interoperability Challenge

Every organization within the infrastructure industries – small to large – faces these interoperability issues at some level. How do we share information among the multitude of application systems across functional boundaries, across organizations, and across the lifecycle? Can the manner in which we deliver information to other participants in the project enterprise add value to our services and products? How do we create and capitalize on new opportunities to leverage information technology and the vast amounts of project information to our collective advantage? How do we capture the knowledge and best practices for project execution so that they can be adapted and reused from project to project?

The following section explores some current approaches to addressing these information technology challenges. This is followed by a description of Bentley's vision for an interoperability platform aimed at removing the current constraints on the effective use of information technology in the context of infrastructure projects. It is this vision that guides Bentley's delivery of an interoperability platform to our users.

> "We too often resort to paper or equivalent 'electronic' paper at significant interfaces, applying manual effort to interpret and input the relevant information into the downstream systems."

Approaches to the Interoperability Challenge

There is a tremendous amount of information created in electronic form within today's infrastructure projects. The previous section discussed multiple issues related to sharing this information and delivering it to other departments and other organizations in a format and scope that is directly usable by the organization receiving the information. There are also issues associated with capturing and reusing the knowledge required to execute an infrastructure project, particularly as it relates to managing the flow of information through the project, which is critical to its success.

We have all seen cases where information in electronic form is in large degree discarded at the interface between two organizations because it was not decipherable or usable by the receiver. There are many potential opportunities for software applications in areas that have received little attention, such as construction, because they are so dependent on receiving information from other organizations, such as design and engineering. We too often resort to paper or equivalent "electronic" paper at significant interfaces, applying manual effort to interpret and input the relevant information into the downstream systems.

The issues associated with the pervasive lack of interoperations among our information technology systems have been around for some time and the "pain" associated with these issues is felt by many. As a result, numerous approaches and strategies to address interoperability have been tried in one form or another. Most of these approaches fall somewhere along the spectrum between the two concepts described below, typically clustering around these two poles.

Wiring Together Islands of Automation

Without a broad strategy that recognizes the value and benefits of interoperability, specific interoperability issues are often identified when they are most painful, such as during the ramp-up or in the middle of executing an infrastructure project. Without an appropriate technology platform to address these issues when they arise, we often resort to the most expedient and tactical approach to fix the immediate problem.

File Export/Import

Perhaps the most pervasive method of achieving interoperability between two applications is by exporting data from one application and importing it into another application, using a format that both applications understand. For software vendors focused on desktop applications alone or not directly concerned with broader interoperability issues, this is often the approach adopted within their applications. This is fine for an ad hoc approach but is certainly no strategic solution to interoperability. Some of the issues with this approach are:

- It is largely a manual process (i.e., running the export and subsequent import) without any managed process to automate the export/import, manage the resulting files, ensure that the process is complete, and react to changes in the source information in a timely manner.
- There isn't a mechanism for change management that automatically identifies the specific changes in the source information and then manages the process of updating the destination application to reflect only those changes. Without such a mechanism, the destination application must reload all of the information from the source application (perhaps overwriting modifications made in the destination application) or the changes must be identified manually and then the destination application updated manually.
- Lastly, the integration is constrained by whatever the common format supports. For example, the format might be incomplete when it is compared to the interoperability requirements. Or the support for the format by the source and destination applications might be incomplete and might well be the lowest common denominator.

"Given the dynamics of infrastructure projects, the exact pairing of source and destination applications might not exist on the next project."

While this approach might have a role in a broader interoperability strategy, it is completely inadequate as an interoperability solution.

Point-to-point interfaces

When file export/import options don't exist or are inadequate, developing specialized point-to-point software interfaces is often employed. But without an integration platform, this approach requires specialized software development to implement. While this approach can potentially address the change management issue specific to the interface between the source and destination application – but only if change management is considered in the implementation plan – it still suffers from the following problems:

- This approach requires software developers to implement the solution. Most infrastructure projects are not staffed with this type of resource and the cost of such an initiative is typically not budgeted.
- Given the dynamics of infrastructure projects, the exact pairing of source and destination applications might not exist on the next project, rendering the integration initiative a "one-off" effort.
- The solution is "fragile" in that it is subject to changes in either the source or destination application. Something as simple as an upgrade to either application can render the point-to-point integration software inoperable. This is certainly true if either application is changed out and replaced with a different application, such as one that provides improved functionality to the end user.

Issues

But, perhaps the single biggest problem with both of these approaches is scalability. As more and more applications are "wired" with either of the above approaches, the magnitude of the above issues grows exponentially as each new application is added. The number of specialized point-to-point integrations could approach $n!$ interfaces, where n is the number of integrated applications. This is simply not tenable for a project of any size. First, the potential of interoperability is constrained by the limitations of the individual interfaces. The point-to-point interfaces ultimately become the Lilliputian ropes that tied Gulliver to the ground. Second, the participants in the project, from an interoperability point of view, are left completely to their own devices – "the proof is left to the reader." The project bears nearly the entire "interoperability tax" on its own.

"Very few infrastructure projects have used a centralized database successfully on a large scale, particularly when considering interoperability across functional and organizational boundaries."

The value of these approaches can never exceed the sum of their parts. While information may move from point to point and a subset of project information might be visible through individual applications, there is no broad, integrated view of the project information. This precludes making strategic use of this information to gain a strategic advantage for the project as a whole or its participants individually. Besides, very little knowledge is "captured" using these approaches. The limited knowledge represented in the point-to-point interfaces is locked into the software and does not represent any process-oriented knowledge.

Centralized Database

Some companies have attempted to solve this problem by using a centralized database. This approach is conceptually straightforward and so easy to understand that it seems to be an obvious choice. In fact, the centralized-database approach is very appropriate within a well-defined scope for processes that remain relatively stable over time. These processes include tightly integrated practitioner workflows and accounting systems. Yet, very few infrastructure projects have used a centralized database successfully on a large scale, particularly when considering interoperability across functional and organizational boundaries. The approach does not satisfy every integration requirement and suffers from a number of issues including data-modeling effort, fragility, vendor and technology lock-in, incompatible transaction model, and scope.

Data-modeling effort

A significant degree of data modeling is required up front. But because all of the data elements within the integration scheme must be identified at this time, incremental implementation becomes very difficult. Implementing a single database often assumes or, depending on the technical implementation, *requires* that all information types, properties, and their relationships are known from the outset. This is particularly true when, as with most of today's implementation, the data model is largely "hard coded" into the database application.

Fragility

The data model is also fragile and inflexible in areas where the information is dynamic, such as newly automated functions – those unique to specific business lines and project types – or opportunities to incorporate new technology.

Vendor and technology lock-in

This approach can leave the user organization dependent on a single database technology from a single vendor for all data requirements. Therefore, it limits the application choice for any given application area. Interactivity with the database using this approach becomes of paramount importance even to the extent that the application fails to deliver on the user needs. Any modification of that application is now interdependent on implementation (and often upgrade) of the database, which causes unnecessary delay in releasing new functionality within the application itself.

Incompatible transaction model

> "A project is a network, not a hierarchy…no one that possesses or requires a 'complete' view of the project. All project participants have their own view of the project, which is 'complete' from their point of view."

As mentioned, the central database approach is entirely appropriate for tightly integrated, well-defined, and focused processes. But when it comes to cross-functional, cross-organizational interoperability, the transaction model is not online database transactions as would be used by a banking system. It usually takes the form of packets of information, often referred to as "work packages," which is an entirely different transaction model. Even within a single organization or workgroup there is very often a requirement to work asynchronously in a disconnected (or "briefcase") mode.

Scope

Lastly, even if the other issues could be successfully addressed, it is inevitable that the data model will never be big enough to encompass all of the data related to a project. Inevitably, there will always be something that falls outside the scope of the database. Information that falls into this category leaves us with the exact same problem we set out to solve – integrating information among dissimilar systems.

Centralized database systems work well for a very well-defined scope. There is, however, a great temptation to let the scope grow. If the centralized database is the only interoperability game in town, everyone wants to make sure that their information requirements are satisfied. The danger is that the scope of the databases can become so large that they collapse under their own weight.

Issues

This leads us to the most significant issue concerning the centralized database approach. It requires a "top-down" implementation. However, with nearly every infrastructure project, there is no "top-down" view. For many, the owner of the asset project is assumed to have such a view, but even the owner answers to its own masters – multiple government regulations and codes, their shareholders, and their customers, to name a few. Further, the owner really can't afford to care about the details of individual, perhaps even unknown, project participants deep in the project supply chain. The fact is, at least from an information point of view, a project is a network, not a hierarchy, so there is no one that possesses or requires a "complete" view of the project. All project participants have their own view of the project, which is "complete" from their point of view.

From a knowledge point of view, the knowledge represented by the information model within the database and its associated applications suffers from the same issues of data modeling, fragility, and technology lock-in as from an interoperability point of view.

Also, as with point-to-point integration, the knowledge is locked away in the associated applications. Finally, from a scope point of view, there is no knowledge captured relative to how the knowledge within the centralized database relates and flows with information beyond the scope of the database.

Bentley's Interoperability Platform Vision

This section outlines Bentley's vision for an interoperability platform that is aimed at transcending the spectrum between the point-to-point and centralized-database approaches. The objective is to outline an approach to interoperability that provides the flexibility, scalability, transparency, and economy demanded by actual infrastructure projects and eliminates the constraints imposed by the other approaches previously discussed. This vision continues to drive Bentley's implementation and support of an interoperability platform on behalf of our users.

"It is the logic of man's situation that prevents him from controlling the drift of history. He can either embrace a social philosophy that fails to see this logic and tries to control the future (futilely, of course), or he can admit his powerlessness with respect to the future – his limits as a mere human – and embrace a social philosophy that upholds those institutions that do in fact leave the future unplanned – and free to make its own mutual adjustments."

–Michael Polanyi, 1969

Approach

It's natural for our rational minds to conclude that if we know the requirements, we can build the system. Putting aside the nontrivial issues associated with "building the system," this seemingly obvious conclusion overlooks an important fact. Given the dynamics, uniqueness, and limited timeframe of each individual infrastructure project, we can *never* know *all* the requirements of a *future* project. So the issue is not one of becoming better predictors of the future. It is not one of capturing every requirement from every point of view in a single place. Rather, it is one of determining how we implement an interoperability strategy that enables us to navigate and thrive in a constantly changing world with constantly changing requirements.

At Bentley, we believe that achieving our objective begins with an interoperability platform. In this section we'll outline the requirements and approach for implementing such a capability.

Requirements for an Interoperable Project Environment

Before defining the specific characteristics of an interoperability platform, it is first useful to examine what the qualities of an entire *interoperable environment* – applications, information, knowledge, and an interoperability platform – should exhibit in order to address the interoperability issues for infrastructure projects. Successfully meeting the interoperability challenge for these projects requires that the interoperable environment exhibit operational independence, managerial independence, geographic distribution, incremental implementation, and emergent behavior.

Operational independence

Each application (or subsystem of applications) within the interoperability environment must be able to operate and provide value in its own right independent of the interoperability strategy and its specific configuration. In the presence of an interoperable environment, each application may and perhaps should perform better and create additional value through information it shares via the interoperable environment.

The same should also be true if the application connects only occasionally to the interoperable environment. However, its basic operation should not be dependent on the existence of the interoperable environment or one specific configuration vs. another. If key applications cannot provide their basic function in the absence of an interoperable environment, the interoperable environment becomes a constraint and the single point of failure. This is a key element required to implement the flexibility and scalability required by infrastructure projects.

Managerial independence

Hand-in-hand with operational independence is managerial independence. Given the multitude of groups and organizations involved in infrastructure projects, a mix that can also vary over time during a single project, the success of an interoperable environment should not be dependent on how the individual applications and sub-systems are acquired and who operates them.

> "In an environment… where we don't know everything up front, it is an absolute necessity that the interoperability platform enables and supports incremental, evolutionary implementation."

Geographic distribution

A critical requirement is the ability to support the increasingly distributed – often globally distributed – nature of infrastructure projects. This means more than assuming that every project participant is connected on the same wide-area network. It means that the interoperable environment can function across multiple communication infrastructures, accommodating whatever communication infrastructure exists in various areas. This means the interoperable environment should accommodate everything from operating in a disconnected mode, to low-bandwidth connections, to the available Internet connectivity, to high-bandwidth corporate networks. It also means extending the "reach" of the project information beyond the desktop to (a) incorporate other technologies for delivering information, such as mobile PCs, PDAs, cell phones, and even advanced uses of paper; and (b) incorporate information from other sources of information, including real-time information through sensor devices such as video cameras, radio-frequency identification (RFID) devices, real-time sensors – including light, temperature, stress, strain, displacement, and environmental – and alarms.

Incremental implementation

In an environment of changing requirements and unforeseen future requirements, where we don't know everything up front, it is an absolute necessity that the interoperability platform enables and supports incremental, evolutionary implementation. Experience, along with new best practices, will continue to inform and, if enabled by the interoperability platform, improve the interoperable environment.

The infrastructure community cannot build a test prototype because the prototype is the project. However, each project could be viewed as a "second-order" prototype from a process and technology point view. As such, each project brings new learning, new practices, and new technology to the next project. However, no standard "product" or "process" is ever realized. Each project becomes a starting point for the next project. An interoperable environment and its interoperability platform must support this type of incremental, evolutionary implementation.

Emergent behavior

In the long term, emergent behavior is perhaps the most strategic requirement. Experience in the infrastructure industries has shown that a federated, system-of-systems approach to interoperability can point the way to new behaviors, new processes, and new opportunities to allow information technology to become more than the sum of its parts by creating new value for the project as a whole and the project participants individually. Of course, this will require that, first, an interoperable environment be implemented with the potential to suggest these emergent behaviors. Second, we must be, collectively, open to new opportunities to change our processes, business models, and deliverables in light of these opportunities. The ability to capitalize on new opportunities *when they present themselves* is a key strategic element of the interoperable project environment.

"Men and nations start with a vague notion of being rich, or great, or good. Each step they make brings unforeseen chances into sight, and shuts out older vistas, and the specifications of the general purpose have to be daily changed. What is reached in the end may be better or worse than what was proposed, but it is always more complex and different."

–William James, 1906

These requirements – operational independence, managerial independence, geographic distribution, incremental implementation, and emergent behavior – are necessary for any significant system to provide value within the context of the uncompromising dynamics and far-flung distribution of infrastructure projects.

Interoperability Platform Characteristics

To overcome point-to-point and centralized-database limitations when applied to infrastructure projects, an interoperability platform should adopt a "federated" approach. Federating project information differs from a centralized-database approach in that the project information is not merged. In a federated approach, multiple heterogeneous applications and information sources are loosely coupled for (a) sharing and synchronizing information between applications, and (b) aggregating information into multiple consolidated "views" of the project information, which are tailored to the requirements of key project participants. In this scenario, each application maintains its autonomy but is able to participate in an interoperable project environment.

In order to create a federated interoperable project environment, the implementation of an *interoperability platform* must include some basic characteristics and functionality. In the following sections these characteristics and functionality are outlined.

Flexible framework

Successful implementation of an interoperability platform must provide a flexible framework for integration that allows for incremental implementation and enables applications to be easily added or replaced. A flexible system makes it possible to adapt the integration scheme to a variety of project types and scopes without having to modify either the application software or the integration software. The key capabilities of this framework include embracing existing applications, extending and modifying dynamic information models, and managing, transforming, and distributing project content.

Embracing existing applications

The interoperability platform should allow existing applications to be *connected* to the interoperable environment while minimizing the impact on the existing applications. The framework should provide configurable connectors for popular industry applications and industry-standard information formats. It should also provide a robust software toolkit for streamlining the implementation of custom or unique connectors. In all cases – unlike the point-to-point approach – only a *single* connection to the interoperable environment should be required.

“Information models that are ‘baked in’ to the source code of the interoperability framework place significant constraint on the flexibility, which is needed to support an incremental and scalable implementation.”

Dynamic information model

Ideally, the interoperability platform should extend and modify the information models without needing subsequent modifications to each of the connections. This would require that the schema for the information model be managed and applied externally to the application software. Information models that are “baked in” to the source code of the interoperability framework place significant constraint on the flexibility, which is needed to support an incremental and scalable implementation.

Managing Project Content

The interoperability platform should include the capability to comprehensively manage the information content critical to a project’s success. Key content management capabilities should encompass the following:

- *All content types* – The interoperability platform should provide the ability to manage all types of project content created by project applications, regardless of whether a Bentley application or a third-party application creates the content. This should include all project content formats including documents, drawings, models, images, and components.
- *Global access* – It is imperative that globally dispersed project teams have access to project content and information. This implies that the content might be effectively managed while physically distributed worldwide. Information access should be provided through thick clients, web clients, and directly within project applications, and should be managed globally.
- *Managing dependencies* – The network of dependencies within and among all the types and representations of project information is incredibly rich and complex. In one sense, these dependencies become the map that shows how a project “works.” From an information technology point of view, some dependencies within narrow domains are represented explicitly within specific applications. Other dependencies might not be explicitly represented but easily inferred. Still others, such as a group of documents that are dependent on a specific paragraph within a contract document, might not be represented and not easily inferred at all as they might exist “in the heads” of key project members.

  The ability to represent, manage, and bring transparency to this map of dependencies within the collective project information – without resorting to the centralized-database approach – gives the project team the ability to quickly determine the impact of proposed changes, evaluate alternatives, validate that project requirements are met, ensure that critical dependencies are not overlooked, and "package" information for specific uses. Therefore, the interoperability platform should provide the capability to (a) display and navigate dependencies explicitly represented within and across applications; (b) create new dependencies within all types of project information; and (c) perform analyses based on the modeled dependencies (for example change-impact analysis).

- *Concurrency control* – Finally, the framework within the interoperability platform should provide capabilities for automatically moving and updating information between applications and information sources in order to maintain concurrency among the multiple sources of information. This "synchronization" may be based on changes in state (such as when a set of drawing is Issued for Construction), project milestones (for instance a 30 percent review), on a regular basis (once a day, every week, every month), or simply on demand. This is different from real-time propagation of every change, or every transaction as it happens, as in a banking system. The *necessarily* asynchronous, highly parallel nature of infrastructure projects, with constant change and exploration of often-conflicting alternatives, would quickly grind to a halt if every change were immediately propagated throughout the entire project.

"The necessarily asynchronous, highly parallel nature of infrastructure projects… would quickly grind to a halt if every change were immediately propagated throughout the entire project."

Transform content

The interoperability platform should provide tools and capabilities for transforming the content from multiple applications in order to (a) pass the information on to other applications connected to the interoperable project environment in order to maintain concurrency as described above, and (b) provide a federated "view" of the project information to project participants. The specific capabilities to achieve this include the capability to:

- *Normalize* the content from multiple applications into one or more information models as dictated by the needs of the destination application or user. As referenced above, it should be possible to easily define, extend, and modify information models – with little or no impact on the individual applications and their interfaces – and then map them into the interoperable environment. To illustrate, multiple applications may represent similar components with entirely different information models. For example, one refers to a door, the other refers to a *wall opening*; the door has a width and the *wall opening* has an *x dimension*. In an integrated, normalized representation, the naming conventions, properties, units of measure, etc., are consistent no matter which source application originally produced the information.
- *Aggregate* the content from one or more applications into higher-level representations or groups, as dictated by the appropriate information models and the level of detail required by the destination application or user. This would include, for example, aggregating multidisciplinary design information. It would also include grouping individual components into a single object to be more directly useable by the destination application or user, such as grouping piping elements into pipe spools.

- *Decompose* the individual components into lower-level components, as dictated by the appropriate information models and the level of detail required by the destination application or user. For example, a single object in a BIM model, such as a wall, may decompose into its individual components (e.g., framing, drywall, wall covering) for use in construction estimating and planning.
- *Filter* the content in order to distribute the appropriate scope of the transform content as required by the destination application or user. For example, if the destination application is used in construction, it might be useful to apply a filter to deliver components from all disciplines within a single area.
- *Reorganize* the information such that it's immediately usable by the destination application or user. For example, engineering may organize by systems covering all areas of the infrastructure asset, whereas construction may want information organized by area, where each area includes all systems. An interoperability platform should provide capabilities that achieve this.

"By managing this knowledge externally to the software under the control of the user, it can be captured, reused, saved as best practices, and quickly adapted to the specifics of the next project."

Distribute content

The interoperability platform should provide robust capabilities to distribute the transformed content in a number of forms, including:

- *Application connectors* – The interoperability platform should distribute project information, once transformed as necessary, to other applications participating in the interoperable environment. It does this through the connectors (as defined above) implemented for those applications. As described above, the interoperability platform should provide software tools to easily implement these connectors where they don't exist.
- *Native file formats* – The interoperability platform should be able to distribute project information in the form of files in the native formats of key industry applications (e.g., DGN, DWG, SHP, PDF, etc.).
- *Industry standards* – Robust industry standards exist in multiple domains within the infrastructure industries (e.g., IFC, OGC WMS, OGC WFS, OGC GML, ISO 15926, LandXML, etc.). The interoperability platform should be able to distribute project information represented in multiple industry-standard forms as applicable.

To provide the flexibility required to effectively support the dynamics of infrastructure projects, it is critically important that the rules for managing, transforming, and distributing project information, as described here, are defined, maintained, and managed *externally* to any compiled software under the *direct control* of the project. These rules, along with the dynamic information model, represent the *knowledge* that manages the information flow and processing for the project. By managing this knowledge externally to the software under the control of the user, it can be captured, reused, saved as best practices, and quickly adapted to the specifics of the next project.

Tools to consume and leverage content

In addition to providing a flexible framework for federating the multiple sources of project information, the interoperability platform should provide a complementary set of client applications, as well as tools for creating new client applications, for consuming federated project information. The functionality of these client applications would include:

- Visualizing federated information (e.g., text, data, drawings, models),
- Querying federated information,
- Commenting and "redlining" federated information,
- Composing project deliverables,
- Supporting new applications for design, analysis, project management, construction, training, etc., which use a federated view of the project information as their starting point.

"An interoperability platform…provides the capability to create.. solutions for specific types of infrastructure assets and solutions for specific lifecycle phases of those infrastructure assets."

These client applications and toolkits should include capabilities for thick clients, web clients, and web portal applications. In addition to the types of applications described above, those provided by the interoperability platform, as well as those created with the software tools, should work in concert with the interoperability platform framework to comprehensively manage the comments, "redlines," etc., created by these applications.

Summary

Beyond enabling an interoperable project environment to achieve actual project interoperations, the interoperability platform described here provides the capability to create true information technology *solutions* – solutions for specific types of infrastructure assets and solutions for specific lifecycle phases of those infrastructure assets. The dynamic information models along with the rules for managing, transforming, and distributing project information represent critical *solution knowledge*, which can be reused and adapted for future projects. These solutions and the associated solution knowledge may be created by Bentley, our users, or our partners. Leveraging the inherent flexibility of the interoperability platform described here allows these solutions to be quickly adapted to the unique realities of each individual project.

Benefits of an Interoperability Platform

The benefits of an interoperability platform might seem obvious, particularly within the context of an individual's own domain and pain points. It would be useful, however, to enumerate just a few benefits from the broad view of the infrastructure project.

Return on information technology investments

As infrastructure projects become increasingly complex and distributed, with greater demands on asset performance and more compressed schedules, information technology becomes an increasingly critical and pervasive tool for successfully managing and executing infrastructure projects. These trends result in increasing investments in information technology. However, it's not just about investing more. More importantly, it's also about realizing a greater *return* on these investments. In this way, an interoperability platform can be an important tool for realizing increased returns.

> "Make better-informed decisions, react more effectively to crises and unexpected events, and explore a broader range of alternatives in project decision-making…these capabilities make it possible to be much more effective in managing risk."

Greater productivity

It's often an article of faith that increased automation results in increased productivity. The broader question is not whether automation can improve the productivity of this task or that task. Rather, it's whether automating a specific task impacts the overall project productivity through, for example, leveraging information created by other applications, providing information for other applications, or providing greater visibility into the state or performance of the project. This is what an interoperability platform enables. As described earlier, the cost to infrastructure projects from the lack of interoperability is significant – conservatively $15.8 billion per year in the U.S. alone. The opportunity is not only to reduce these costs but to further streamline the project processes and execution by identifying new connections and new applications for the project information.

Further, the inherent flexibility of the interoperability platform envisioned here enables projects to adopt the most productive applications for individual tasks and not be constrained to specific applications from a single vendor simply for the sake of compatibility.

Increased flexibility and resilience

A flexible interoperability platform leads to a flexible organization. It makes it possible to adapt to new conditions – new technology, new project types, new business opportunities. It also allows a project to adapt in real time to the inevitable unplanned and unexpected events arising during the project. This flexibility also dramatically improves scalability. The ability to configure familiar and productive tools for small projects, large projects, and diverse projects can be a significant advantage. The flexibility and federated approach described here also enables greater resilience for the project. There is no "single point of failure." In fact, there is no reliance on one single technology or one single software vendor.

Increased visibility into project information content

The federated approach with the ability to create multiple aggregated views into the project information provides much greater visibility into the project and how it is performing. This visibility is further enhanced by the ability to create, visualize, and manage the multiple dependencies within the project information. These capabilities make it possible to make better-informed decisions, react more effectively to crises and unexpected events, and explore a broader range of alternatives in project decision-making. In the end, these capabilities make it possible to be much more effective in managing risk.

Better positioned to identify, create, and leverage new opportunities

"Thinking makes human life possible because it allows man to imagine the future and confront it. If suddenly all interpretations of the world were excised and we were left with no knowledge of what might happen tomorrow – not even the vaguest notion – the terror would overwhelm us. Thus thought makes tomorrow possible, and tomorrow is time itself. Far from being merely a way of passing time, thought creates time."

–Jose Ortega y Gasset, 1940

The strategic benefit of an interoperability platform will ultimately enable project participants to identify and adopt new methods of project execution and new project business models. It directly supports multiple trends among the infrastructure communities – design-build, public-private partnerships, and integrated project delivery, to name a few. It will enable individual companies to operate in new geographic areas or leverage available resources on a global basis. It will also enable these companies to create new products, new services, and new forms of intellectual property.

Capture and reuse critical knowledge

As described earlier, the rules for processing project information, along with the dynamic information models, represent the *knowledge* that manages the information flow and processing for the project. Through applying the interoperability platform to achieve interoperations, this knowledge becomes an *explicit*, and thus *enduring* and *reusable*, representation of best practices for executing projects for specific infrastructure asset classes and specific lifecycle phases. This knowledge can then be immediately reused and adapted as necessary for the next project without relying solely on the experience of past project team members or requiring time for new team members to get up to speed with project procedures and practices communicated in written or verbal form.

Summary

The interoperability platform envisioned here and being delivered by Bentley can be a strategic component of any infrastructure project by creating an interoperable project environment and enabling true interoperations among the diverse software applications applied to the project. This platform delivers the flexibility, scalability, and robustness demanded by today's infrastructure projects. It acknowledges the fact the every project is unique and every project a special case. It is the key to navigating and thriving in a dynamic industry with ever-evolving requirements and opportunities arising as we collectively continue to create, operate, and sustain the world's infrastructure.