Enterprise Search in SharePoint 2010

The SharePoint 2010 lineup includes improvements to its enterprise search capabilities, particularly in the areas of relevance, user refinement of search queries, people search, and scalability. For customers without an enterprise search solution, the improvements could bolster the case for deploying SharePoint Server 2010, and smaller businesses could benefit from using the low-cost Search Server or free Search Server Express. Organizations that need more robust enterprise search capabilities and scale should consider the new (and more expensive) FAST Search Server 2010 for SharePoint.

Enterprise Search and SharePoint

Enterprise search tools help workers locate internal information in a wide range of formats stored in intranets, messaging systems, shared files on networks, and databases. Enterprise search products use components called *crawlers*, which automatically examine information sources on networks to retrieve data stored within files and databases and then create *indexes* that contain pointers to information discovered. Users conduct search queries by entering keywords into a Web-based interface. These queries are run against the index, and the system returns results ranked in order of relevance. The user experience and underlying technical components are similar to those of Web-based search engines such as Google or Bing, but at smaller scale, and with support for indexing more types of data.

Microsoft has built search capabilities into various products over the years, but around 2005 it began concentrating its enterprise search efforts on SharePoint Server. (For an overview of other Microsoft products with search capabilities, see the sidebar "Search in Other Microsoft Products".) Bundling enterprise search into SharePoint let Microsoft offer a much cheaper (if less functional) alternative to high-end enterprise search products from providers like Autonomy, Endeca, and FAST (later purchased by Microsoft), and had technical benefits as well, such as helping users find other employees with particular types of expertise. SharePoint's enterprise search features also provided a counter to Google's enterprise search product, the Google Search Appliance, which can index internal resources including SharePoint sites, and which gave Google an entryway into enterprises.

In 2008, Microsoft's enterprise search strategy was affected by three main factors. First, the company updated SharePoint Server 2007 and a search-oriented subproduct (renamed Search Server 2008) with several important improvements, including *federated search*, which is the ability to merge results from external search engines, such as those on partner extranets and external Web sites. Second, the company released a free version, Search Server Express 2008, for small businesses that need a basic one-server solution. Third, Microsoft acquired Norway-based FAST Search and Transfer in 2008 and began selling FAST ESP (based on the acronym for "enterprise search platform"), which could index far greater quantities and more types of information and handle more simultaneous queries than SharePoint, and gave enterprises more control over how results were ranked.

Absorbing FAST

With the release of the SharePoint 2010 lineup in May 2010, Microsoft will take its first steps toward integrating FAST technology into SharePoint.

SharePoint, Search Server, and Search Server Express all get new releases in 2010, and all improve search in areas such as relevance and user refinement of search queries. These improvements are not based on FAST technology.

Microsoft is also introducing a new high-end product, FAST Search Server 2010 for SharePoint, which uses SharePoint Server 2010 as the front end for user queries and some types of searches (such as people search), while relying on FAST for core search functionality, such as indexing documents and ranking results. Microsoft recommends this product for companies that want to better "tune" search results to improve their relevance or meet unique business needs, or to build customized internal search applications. FAST for SharePoint can also scale up to support larger index sizes and provide better query performance than SharePoint. (For a chart showing a basic comparison of the four SharePoint enterprise search products, see "SharePoint Enterprise Search Products Compared".)

At its SharePoint conference in Nov. 2009, Microsoft said it plans to build a single search platform for all its products, suggesting that the technical split between SharePoint and FAST will gradually disappear. Microsoft has also begun to pare back FAST offerings in nonstrategic areas, selling some former FAST products and discontinuing others. Most notably, the company said it would stop developing FAST ESP for Unix after 2010. This suggests that SharePoint remains the main focus of Microsoft's enterprise search product line, and FAST technology will primarily be used to improve SharePoint's search capabilities.

Enterprise Search Improvements in SharePoint Server 2010

SharePoint Server 2010 includes numerous enterprise search improvements over its predecessor in areas such as relevance and core search, people search, and extensibility, and will benefit from a new scale-out architecture that will let it index more documents and handle more simultaneous queries.

Search Server 2010 and Search Server 2010 Express also get most of the new features, although enterprise search features that rely on other SharePoint Server 2010 features won't work in these lesser products, and there are significant differences in scalability as well.

With their 2010 versions, Search Server and Express will for the first time be able to use SharePoint's Business Connectivity Services (BCS) connector to index data from line-of-business applications like those from SAP and Siebel.

Relevance and Core Search

SharePoint 2010 improves search results in two main ways: by letting users enter and refine their queries more effectively and through tweaks to the algorithms used to display results in order of relevance.

Many query improvements add features that have long been available in Web and high-end enterprise search engines. For example, users can now enter Boolean (e.g., "AND," "OR," "NOT") and wildcard ("*") searches and search on dates or numerical ranges by using symbols (e.g., "<", ">", and "="). In addition, SharePoint 2010 will autocomplete suggested query terms, based on past searches by other users.

Once users receive results, SharePoint can suggest related searches based on similar terms searched by other users. A new refinement panel lets users filter results by content type, location, author, date, or *managed properties*, which are descriptive metadata terms for documents defined either by their creators or by organization-wide rules, such as product taxonomies. (Because managed properties requires SharePoint Server, this feature is not supported on Search Server.)

SharePoint 2010 can also infer the title and author from data within the document itself (e.g., a line "By Matt Rosoff" on the first slide of a PowerPoint presentation) and then consider this information when ranking results.

Improvements to SharePoint's relevance-ranking algorithm take into account more factors, such as the number of times users have clicked through to a particular document, "social tags" (a new feature that lets users tag resources with particular keywords and rank their relevance or usefulness), and "fuzzy matching" for URLs (wherein a URL that contains a term similar—but not identical—to the search term will show up in results). SharePoint 2010 also does a better job of understanding and identifying multiple languages and breaking words accordingly (e.g., compound words in German can be divided into component parts).

People Search

Microsoft has promoted certain SharePoint Server features as "corporate social networking." In particular, the product lets employees create personal home pages known as My Sites. These pages show up in search results, allowing users to find others in their organization with particular areas of expertise.

SharePoint Server 2010 offers a number of improvements in how users can find other people in their organization, such as the following:

- Queries on nicknames ("Matt" instead of "Matthew") and phonetic misspellings ("Roseoff" instead of "Rosoff") are understood (this requires installing the optional Microsoft Speech Platform, available through SharePoint's software prerequisites tool)
- A person's areas of expertise can be inferred from keywords in the e-mail he sends and receives (an Outlook 2010 add-in is required)

- Relevance ranking now takes into account "social distance," or the organizational distance between the searcher and the individual in each result based on organizational hierarchies in Active Directory, as well as the number of times the searcher has clicked on a result in the past.

Organizations should note that these people-search results are only as good as the data in the user's profile—for instance, users may misrepresent or overestimate their areas of expertise in their My Sites profiles.

Because My Sites and other social-networking features are not included in Search Server, people-search features require SharePoint Server 2010. (FAST Search for SharePoint also uses SharePoint Server for people-search features, while relying on the FAST back end for other types of enterprise search.)

Scale-Out Architecture

A new architecture for enterprise search in SharePoint Server 2010 lets organizations add more servers to reduce latency time for queries, improve the freshness of indexes, and provide greater resilience and availability. These improvements also allow a properly scaled SharePoint 2010 Server installation to index up to 100 million items—twice as many as its predecessor.

SharePoint 2007 enterprise search installations included one or more indexing servers, each of which crawled all resources and built a complete (and redundant, in the case of multiple servers) index. Although organizations could improve performance by adding servers, index size was limited and very large indexes could cause slow queries. Moreover, if an indexing server became too busy or went down, it affected both query time and crawling performance.

With SharePoint 2010, a SharePoint installation can have multiple crawlers indexing different resources simultaneously, improving crawl performance and the "freshness" of the index. The resulting index can be partitioned and mirrored, and associated crawl and property databases can be clustered and mirrored for performance and availability. (For an architectural diagram and more details, see the illustration "Enterprise Search Architecture in SharePoint 2010".)

Search Server 2010 can be scaled in the same way, but Search Server 2010 Express is limited to a single server, and therefore cannot. Microsoft says that Search Server 2010 Express can index up to 300,000 items if running in conjunction with SQL Server Express, and up to 10 million if using a full version of SQL Server.

Extensibility

SharePoint 2010 enterprise search can be extended or customized in new ways.

Search Web Parts let administrators or power users create customized search experiences, such as predefined queries, a Best Bets panel (which highlights suggested top results for particular queries), featured content, and automatic filtering by category

(e.g., date, author, managed properties). Unlike previous versions, these Web Parts can be modified by developers; prior search Web Parts were "sealed," which prevented changes.

Although SharePoint 2007 lets users access search results from external providers (federated search results), developers could not access federated data except to display results. In SharePoint 2010, a new federation object model lets developers build and extend Web Parts that run specific queries against one or more external sources or *search providers*—such as other SharePoint and FAST search indexes, application databases, and search engines on external Web sites—and present results in a single interface. Developers can also extend this object model to query other types of external resources.

Finally, administrators can connect SharePoint 2010 to Windows 7 Search. This allows users to save specific SharePoint queries and access them later from the Windows 7 Search pane, without opening a browser and navigating to the SharePoint search interface. This feature does not enable SharePoint Server to search resources on a user's desktop but provides a more convenient way for a user to search using SharePoint Server.

FAST Search Server 2010 for SharePoint

FAST Search Server 2010 for SharePoint (FAST for SharePoint) uses SharePoint Server 2010 to provide user and administrative interfaces and some other features, while a new FAST server provides most core search functions, such as crawling, indexing, and relevance ranking. (For an illustration showing which features are provided by which components, see "FAST for SharePoint Architecture".)

FAST for SharePoint is intended for customers who plan to deploy SharePoint Server 2010 but want better enterprise search functionality, such as the following:

- More granular controls for users, administrators, and developers for tuning and customizing search results
- Indexing more than 100 million documents
- Improved query performance by scaling out—useful for large numbers of simultaneous users.

Microsoft will continue to sell a stand-alone version of FAST for enterprise search, called FAST Search for Internal Applications (formerly FAST ESP). However, the integrated FAST for SharePoint product gains a number of benefits from SharePoint, such as the following:

- People-search functionality—these results are provided by SharePoint, not the FAST components
- An extended version of the SharePoint Search Center interface, a Web-based interface for queries, which will be familiar to SharePoint users
- Administrators can use SharePoint's familiar Site Settings and Central Administration interfaces to fine-tune results and implement certain features (such as extracting specific properties from documents to use in relevance rankings)

- PowerShell components for administrative scripting.

FAST for SharePoint will also be the first FAST product available on volume licensing price lists and sold through Microsoft's traditional partner channel. Previously, customers could purchase FAST only through the dedicated FAST partner channel. This FAST channel will continue to sell FAST Search for Internal Applications.

Relevance and Core Search

Compared with SharePoint 2010, FAST for SharePoint can gather more data from within documents to ensure that search results are arranged in order of relevance. *Property extraction* (sometimes called entity extraction) extracts data such as location, people, and company names from the body of an indexed document and uses that information when determining relevance. Administrators can include or exclude specific properties and can create custom property extractions based on their own content (e.g., the name of a particularly important project could be extracted).

FAST for SharePoint also offers more granular tools for tuning search results for relevance and customizing results for specific users and groups. For example, administrators can do the following:

- Set the relative weights of properties like freshness (how recently the document was added or updated), proximity (distance between query terms), authority (external links to a document—similar to Google's PageRank system), query authority (the number of other users who have clicked through to a particular result), and specific values of SharePoint managed properties (e.g., weigh documents by authors in the "executive" group more heavily than documents by authors in the "copyeditor" group)
- Promote or demote search results for specific documents to move them up or down in the relevance rankings and customize these results for particular groups of users (e.g., so a user in the "editorial" group is more likely to see a style guide document than a user in the "sales" group)
- Finally, the product can index more types of data than SharePoint 2010. For example, it offers broader language coverage, detecting 84 languages and applying "advanced linguistics," such as word breaking, synonym identification, and lemmatization (e.g., identifying a term like "better" as being related to its lemma, "good"). In addition, the Web site crawler for FAST for SharePoint can index more complicated Web pages by pulling down both HTML and associated files, such as JavaScript scripts and Cascading Style Sheets (CSS).

Users Get More Control, Visual Results

In addition to providing tools for administrators to tune relevance, FAST for SharePoint also provides more ways for users to tune and narrow search results, such as the following:

- Users can sort results on a wide range of properties, such as author, document, or title; SharePoint 2010 results can be sorted only by relevance and date

- A deep refinement pane, which suggests ways to narrow searches to get fewer (but more relevant) results, includes the precise number of documents in a wide range of categories (e.g., 14 Word documents and 3 PDFs; 5 by author Matt Rosoff and 12 by author Rob Helm) drawn from SharePoint metadata, property extraction, or administrator-set criteria based on SharePoint-managed properties; users can then drill down into each subset of results
- A "similar results" link appears below some results, letting users conduct a secondary search limited to similar documents
- Duplicate files that appear in multiple places will be collapsed into a single result
- Sophisticated users can use the FAST Query Language (FQL) to create, save, and edit customized search queries—for example, to limit matches to particular fields (e.g., title, body) and to set the weight of particular properties using the XRANK tool (e.g., to ensure that documents with the words "Matt Rosoff" in them appear higher in search results)
- In addition, although the search interface looks similar to that of SharePoint 2010, FAST for SharePoint can include more graphical content in its search results, including Visual Best Bets (in which an image is included as part of a preselected Best Bet for a particular query), thumbnail previews of Word documents, and PowerPoint previews that enable users to scroll through thumbnails of a deck's slides.

Scalability

FAST for SharePoint has multiple components that can be scaled out to improve performance in specific areas.

The item processing service, which crawls data, can be scaled out to multiple load-balanced servers to improve crawling performance (which helps keep results up-to-date), as well as for fault tolerance.

The indexing service, which arranges crawled data into the indexes against which searches are conducted, can be split across multiple indexing columns on different FAST servers, with each column containing a portion of the entire index. This increases the maximum size of the index and speeds indexing time, which helps keep search results up-to-date as the underlying sources change. The entire index can also be backed up in multiple indexing rows for fault tolerance.

The query matching service, which compares user queries against the indexes and returns appropriate results, can be scaled across multiple FAST servers to improve query performance and provide fault tolerance. Here, multiple instances of the query service are deployed as fully redundant search rows that intersect with all the indexing columns. Queries are matched against all the columns independently and then combined in the search results.

Microsoft says that an installation of FAST for SharePoint scaled in this fashion can index more than 500 million documents. The company has not yet published query performance benchmarks.

Extensibility

FAST for SharePoint uses the same object models and APIs as SharePoint 2010, so developers can build one application that will work on both types of system. However, FAST for SharePoint has additional features that developers can tap into, such as the following:

- The customizable search Web Parts include features specific to FAST for SharePoint, such as sorting by managed properties, deep refinements, and similar-searches links
- An extensible content pipeline lets developers program new ways to analyze text within crawled resources—for example, developers could write code to provide in-line translations or classify material into custom types, such as "company confidential" or "expires in 2010."

What to Watch For

Enterprise search installations are not trivial to set up, and partner expertise will likely be required. This is particularly true in the case of FAST for SharePoint, which offers many more ways to tune results and which introduces many components that will be new to current SharePoint administrators. Consequently, licensing and hardware costs are likely to be a relatively small part of the overall expenditure required to deploy SharePoint enterprise search.

Upgrading from SharePoint 2007 could be complicated, since some enterprise search components require SharePoint Server 2010 and won't coexist with earlier versions. Also, SharePoint 2007 was available in a 32-bit version, while SharePoint 2010 is 64-bit only, and therefore will require 64-bit versions of Windows Server 2008 (or later) and SQL Server 2005 (or later).

In the case of FAST for SharePoint, mixed environments are possible: customers will need SharePoint 2010 servers for administration and for certain FAST for SharePoint components, but a FAST for SharePoint farm can index content stored on SharePoint 2007 farms. Customers could also buy FAST for Internal Applications (which doesn't use SharePoint but can index material on both SharePoint 2007 and SharePoint 2010), and then later upgrade to FAST for SharePoint when their SharePoint 2010 deployments are complete. However, this would be a rip-and-replace operation, as the FAST for Internet Applications component cannot be used as the back end in FAST for SharePoint.

Also, some of the feature differences between SharePoint 2010 and FAST for SharePoint seem arbitrary. For example, Microsoft added support for Boolean operators and wildcard searches (which were already available in FAST ESP) to SharePoint 2010, but reserved other features such as Visual Best Bets, similar results, and consolidated duplicates for FAST for SharePoint. In the long run, FAST for SharePoint will probably get more improvements in enterprise search than its lower-priced cousins.

Finally, customers who are considering deploying a non-SharePoint version of FAST should realize that Microsoft will probably focus development on FAST for SharePoint, and that other FAST products might receive less development and less frequent updates over the next decade.

SharePoint 2010 Enhances Content Management

SharePoint Server 2010 offers many improvements for organizing, searching, and archiving documents and other content. It will also help organizations get control of policies and procedures for managing content in SharePoint sites, which have spread rapidly and at times without enterprise oversight. However, SharePoint Server cannot manage content in other stores (such as Windows file systems), and to avoid SharePoint "sprawl" organizations will need to implement effective content management policies and procedures both within SharePoint and outside of it, a task that offers opportunities for Microsoft partners.

Latest Chapter in SharePoint Content Management

Enterprise content management (ECM) is a broad term for processes that manage documents, Web pages, images, and other digital information inside organizations. It subsumes several industry terms, including the following:

Document management systems organize, categorize, and store working documents and make those documents easier for users to find and use.

Records management systems supervise protection and destruction of an organization's records, which consist of documents and other types of data (for example, personal calendars) preserved for legal, regulatory, or strategic reasons..

Web content management systems help users format and publish HTML pages and other types of Web content to sites and ensure that content review and other procedures are followed.

In general, content management systems go beyond simple file storage: they improve search and navigation and capture version histories so that users can recover from errors and identify which authors changed which data and when they did it. They also capture descriptive properties—metadata—for content to aid search and enforce policy (for example, identifying documents containing company-confidential information). Content management systems also provide workflow systems and other mechanisms to help users follow prescribed policies, such as ensuring that all information published to an organization's public Web site has received the necessary approvals and that records of important business events (such as contract signings) are retained for the minimum required periods.

SharePoint Server and the free SharePoint Foundation (formerly Windows SharePoint Services) both include basic content management capabilities such as versioning and workflow for content stored in them. (For a review of SharePoint content organization concepts, see the sidebar "How SharePoint Content Is Organized".) However, SharePoint Server gives

organizations more central control over those capabilities, can scale to large installations, and addresses records management and other needs that SharePoint Foundation does not. Organizations that want to manage SharePoint content according to centrally defined policies and procedures should deploy the full SharePoint Server.

SharePoint Server 2010 continues and extends the content management strategy of previous versions. In general, this version makes content management features more visible and usable so that users are more likely to participate. It also makes policy definition and enforcement more consistent across the product and enables more central configuration so that organizations have more confidence that SharePoint content is under control.

This chapter covers the general-purpose content management capabilities of SharePoint Server 2010, which apply to all types of documents as well as other content, such as calendars and tasks lists. SharePoint Server 2010 also has some specific new features that apply only to particular types of content, such as Office documents, HTML pages and other Web content, and digital assets such as video; these features are covered in the chapter "Web Content Management Improved in SharePoint 2010".

Standardizing Metadata

A new managed metadata service in SharePoint Server 2010 can centrally define and maintain metadata used to organize content on all SharePoint Server sites in an organization. An organization can, for instance, define a standard set of document properties and a dictionary of standard search keywords and mandate them for all SharePoint sites. Centrally defined metadata promote consistency in filing SharePoint content, aid search and navigation, and help enforce policies such as rules for retaining content.

The managed metadata service enables organizations to centrally distribute the following types of metadata definitions to all sites:

- Standard definitions of types of content and management policies for each type
- Standard vocabularies of keywords to aid search and navigation.

Syndication to Standardize Content Types

In earlier SharePoint versions, content types are defined for each SharePoint site collection. However, large, geographically distributed organizations often have multiple site collections or farms. This makes it hard to maintain a consistent set of content types across the organization to promote standard handling for content, aid search, and enforce policy.

The SharePoint 2010 managed metadata service tackles this problem. It can publish a set of content types and associated columns, information management policies and other information to a set of SharePoint installations (farms) and all site collections in those farms. For example, an organization could publish a standard employee performance review document type

with a prescribed set of columns (such as "employee id" and "review period"), a Word document template, and a retention policy, which would be available across all its SharePoint farms.

SharePoint content types can also have associated workflows: for example, an employee performance review type might have workflows to gather reviewer input and to obtain approval for archiving. For each content type, the managed metadata service publishes the names of its associated workflows. However, the service does not actually install associated workflows on each site that uses a content type. SharePoint's built-in workflows install with the product and so will typically already exist on most SharePoint sites. However, custom workflows will not; if a content type has an associated custom workflow, an organization will have to use scripting or another solution to install that workflow anywhere the content type is used.

Taxonomies Aid Search and Navigation

With SharePoint Server 2010, organizations can define hierarchical dictionaries of terms, called *taxonomies*, to use in SharePoint columns; the managed metadata service can then distribute these taxonomies to all SharePoint farms. For example, an organization could define a standard set of terms for its own activities (legal, human resources, manufacturing) and for classifying sensitive information (company confidential, attorney work product). By defining standard terms for document properties and other columns, an organization can improve SharePoint search and navigation and promote consistent handling for sensitive content.

Organizations can define taxonomies using a Web-based administration tool, or import them using a SharePoint-specific text file format. Terms can include translations in multiple languages (useful in global organizations) and can also have defined synonyms (useful if a company isn't entirely consistent in its terminology). Once a taxonomy has been defined, two SharePoint features enable an organization to incorporate it into content.

Managed metadata columns. A managed metadata column is a SharePoint column that takes its values from a defined set of terms in a taxonomy: for example, a "classification" column for documents might take its values from a taxonomy with terms such as "secret," "confidential," "sensitive but unclassified," and "unclassified." Once a managed metadata column has been defined, it can be indexed for full-text search, navigation, and filtering of document libraries and other SharePoint lists. (For specific examples, see the illustration "Navigating Documents with Taxonomies".) Managed metadata columns resemble SharePoint "choice" columns, which also take their values from a defined set of terms, but the terms for a choice column are part of that column's definition. Terms for a metadata column, in contrast, can be shared with other columns and maintained centrally for all those columns, ensuring consistency.

The Managed Keywords column. SharePoint Server 2010 also predefines a column called Managed Keywords, which exploits taxonomies. This column enables users to enter keywords to describe a document or other content to aid search. Users can enter any keywords they like into the Managed Keywords column. However, during data entry SharePoint will automatically suggest taxonomy terms that match what the user is typing. This can speed up keyword entry and make it more consistent, which makes search more accurate.

Centralization Possible, but Not Mandated

The managed metadata service is an option, not a requirement, and multiple instances can run at the same time. An organization might run multiple instances of the service for security reasons: for example, a company's legal department might want to set up its own taxonomy with terms about pending cases that should not be visible outside the department.

An organization that does want a single, centralized set of content types and taxonomies will have to ensure that every SharePoint farm subscribes to the managed metadata service. Among other things, that means preventing independent installs of the free SharePoint Foundation (formerly called Windows SharePoint Services), which cannot connect to the service.

Managing Working Documents

SharePoint Server 2010 introduces new or updated features to keep track of active documents and ensure that they go through an organization's defined approval and publication processes. In general, these features will help users locate documents in SharePoint Server and run documents through prescribed approval or other processes more consistently. Also, both SharePoint Server 2010 and the free SharePoint Foundation 2010 benefit from a new user interface and performance fixes that particularly affect content management.

Filing and Locating Documents

New features for organizing and locating documents in SharePoint sites are displayed in the updated Document Center, a site template for large-scale document stores, but the features can be set up in any SharePoint Server 2010 site.

Document IDs. All documents can be assigned a unique, persistent document ID column that remains attached to the document through successive versions and follows the document when it is moved within SharePoint (although SharePoint assigns a new document ID when it is copied). Users can locate the document by ID within a site collection through a special URL redirector component and across all site collections using SharePoint search. This feature can help users find documents that have been moved, whether by mistake or because they have been placed in an archival location. Organizations can define their own document ID formats (for example, giving document IDs a meaningful prefix) and code their own redirector components (to look up document IDs in an external database, for example).

Document sets. SharePoint Server 2010 introduces document sets, a new kind of SharePoint folder that captures all documents associated with a project, useful for managing the documents and the project with a consistent set of rules. Organizations can define document-set content types as templates for particular projects. For example, an organization could define a document set content type for sales proposals. The type could include template documents (for example, a standard Word proposal document and Excel bill of materials for a sales proposal), a welcome page that appears when an instance of the document set type is opened, and a workflow for management review of the proposal. Document-set content types can define columns whose values will be shared by all documents in the set: for example, a sales proposal document set might have a

"customer number" column. Document sets can have a version history separate from the documents inside them; this enables users to go back to an earlier point in time and see all the project's documents as they were at that time.

Earlier versions of SharePoint enable organizations to define folder content types, but these do not have a starting set of template documents or a welcome page, and associating a workflow with a folder content type requires custom code. Documents in an ordinary folder also are versioned independently of one another; there is no built-in way to capture a consistent snapshot of an entire folder at a point in time and then go back to it later.

Content Organizer. SharePoint 2010 renames and improves the SharePoint 2007 Record Router, an automatic filing component for documents. The Content Organizer automatically moves documents from a special "drop-off library" to document libraries or other destinations in a SharePoint installation, based on a set of routing rules defined in code or in a Web-based rule definition dialog box. Like its predecessor, the Content Organizer can automatically archive documents for longer-term storage, but it can also be used to help users unfamiliar with a site put working documents in the correct place. It also enables scripts or other automated processes to add content to a SharePoint site without having to incorporate knowledge of the site structure, since that structure can be defined in and called from Content Organizer.

The Content Organizer enables users to define more sophisticated document routing rules than in past versions. Among other things, it can route documents based on values in a document's columns: An organization could, for instance, route sales proposals to destination folders based on a customer number column. This contrasts with the earlier version, which routed strictly by content type, and so it would (for example) send all sales proposals to the same destination. The Content Organizer can automatically create new destination folders and can also move a document but leave a link, enabling users to access a moved document from its old location.

Folder metadata inheritance. A SharePoint Server 2010 folder can define default values for columns and information management policies that will apply by default to all content added to the folder and subfolders. For example, a department might specify that all content placed in a particular folder should have its "department" column set to the correct value by default. This feature can help organize content by ensuring that columns are set to reasonable values even if users forget. Developers could code this kind of metadata inheritance in earlier SharePoint versions, but the capability is now built in.

Platform Enhances User Interface, Scalability, Workflow

SharePoint Server 2010 content management also benefits from improvements to the base SharePoint platform, improvements that also apply to the free SharePoint Foundation 2010. Notable improvements include the following:

- A new browser interface that incorporates the Ribbon, a context-sensitive graphical toolbar similar to that of Office 2007 and Office 2010, which can help make content management commands (such as checking out documents and editing properties) more visible to users

- Performance and reliability improvements for SharePoint lists, such as a new "query throttling" feature, can cut off attempts to retrieve more than 5,000 list items at once, which might happen when a user navigates to the default view of a large list or document library
- Workflow improvements, such as the ability to run a SharePoint workflow as a different user, which lets (for example) a user archive documents in a location the user could not write to under her own identity.

Note that the workflow system in SharePoint Server 2010 and SharePoint Foundation 2010 is based on the Windows Workflow Foundation (WF) 3.5 engine, which is largely backward-compatible with the engine in SharePoint Server 2007 and Windows SharePoint Services 3.0. It is not based on the upcoming WF 4.0 engine, which makes major performance and other improvements but is not fully backward-compatible.

Records Management Everywhere in SharePoint

SharePoint Server 2010 has many improvements to help organizations manage records in the product more consistently. In records management, *records* are documents or data that provide reliable evidence of important business events, such the signing of a contract or a decision to make an acquisition. Generally, records need to be protected from later changes and retained for minimum periods of time. Records management is critical to prepare for future legal disputes and to comply with regulations, but organizations also manage records to protect critical documents (such as design specifications and policy manuals) for reference, training, or historical purposes.

SharePoint Server 2010, like SharePoint Server 2007, ships with a Records Center site template for managing SharePoint content as records. In SharePoint Server 2007, records management features were available only in the Records Center; in SharePoint Server 2010, these features can be used anywhere. SharePoint Server 2010 incorporates many other improvements, some of which were originally introduced in a 2008 Records Management Resource Kit for partners. In general, the changes enable records management outside the Records Center and give finer control over final disposition (archiving or deletion) of records stored in SharePoint.

Declaring Records and Legal Holds in Place

SharePoint Server 2010 enables organizations to flag any SharePoint list item or document as a record or subject it to a legal hold to protect it during legal discovery processes. This contrasts with SharePoint Server 2007, whose records management and legal hold features only applied to content in its Records Center site. The change could ensure that records management policies are applied to relevant SharePoint content and simplify records management processes and regulatory compliance.

(For a view of how these features are integrated into SharePoint sites, see the illustration "SharePoint 2010 Records Management and Retention".)

Declaring records. Any SharePoint list item or document can be declared a record and thereby get an information management policy that is different from nonrecords. Organizations can define information management policies that automatically define all content in a list to be records after a given period of time, or create workflows that run on items and declare them records; for example, an organization might create an approval workflow for reviewing potential records in its document libraries.

Placing legal holds. Users with appropriate permissions can search documents and list items from one or more SharePoint sites and place or release a legal hold on items found. Placing a legal hold suspends normal retention policy and prevents modification or deletion. Users can also place or release a legal hold on individual SharePoint list items or documents. These features can help comply with legal discovery requests by locating and locking down content that might be relevant to the request. Organizations can also generate a report in Excel that summarizes the set of items currently on legal hold.

Disposing of Content

SharePoint Server features such as information management policies help dispose of old content, which improves search and navigation, simplifies future legal discovery requests by eliminating irrelevant material, and saves disk space to simplify backup and restore. In SharePoint Server 2010, these features have been improved to handle the more complex disposal processes commonly required by records management.

Multistage disposition. SharePoint Server 2010 can define multistage disposition processes based on the length of time an item has existed or has been unchanged. For example, a department might declare all items in a particular document library as records after a year, review them annually, and delete them permanently after five years. In earlier versions of SharePoint, organizations had to create custom workflows or code to implement multistage disposition.

Some disposition processes still require custom workflows or code in SharePoint Server 2010. For example, many organizations dispose of records based on specific business events rather than elapsed time: An organization might delete or archive employee records based on the employee's termination date, or contracts based on the contract term. Such processes will still require custom workflows and/or code in SharePoint Server 2010.

Hierarchical file plans. SharePoint Server 2010 includes tools for managing records according to hierarchical file plans—a critical tool for managing records in many large organizations, particularly in the public sector.

A hierarchical file plan defines a hierarchical classification scheme for records, with retention and disposition rules for each class. In SharePoint Server 2010, organizations set up hierarchical file plans by creating a folder structure, with a folder for each class of records that defines the policy for that class. Each folder inherits information management policy from its container by default but can override the inherited policy with its own. (This is similar to the way access control lists work in Windows file folders.) Administrators can generate reports to get an overview of the hierarchy, such as the number of records in each folder.

Word Automation Services. SharePoint Server 2010 includes a new developer component called Word Automation Services that can convert Word documents into other formats, which could be particularly valuable for public sector organizations with specific long-term archiving requirements. (For a summary of the component, see the sidebar "Word Automation Services for Document Conversion".)

Records Beyond SharePoint

The new records management features in SharePoint 2010 help organizations manage SharePoint content consistently, but do not aid management of other data sources, such as Windows file shares and Exchange mailbox databases. Furthermore, records management policies in SharePoint cannot be synchronized with policies outside it; for example, Exchange has its own retention policy and legal hold mechanisms, which do not coordinate with those of SharePoint. In general, all material to be managed as records by SharePoint must be brought into SharePoint.

Some Microsoft products can contribute content to SharePoint Server, which enables organizations to use SharePoint as their central records store. Exchange Server 2007 and 2010 have a Managed Folders feature which, among other things, can mail copies of messages to another location; that location could be a SharePoint list. Similarly, organizations can use the Windows Server 2008 R2 File Classification Infrastructure to scan and automatically add properties to files; when a file with classification properties is uploaded into a SharePoint document library, SharePoint will preserve properties that match corresponding SharePoint columns, provided that the file is an Office document or some other file type known to SharePoint.

Making Content Management More Accessible, Consistent

Overall, the new capabilities of SharePoint Server 2010 could make content management features simpler, which in turn could encourage users to cooperate with content management policies and procedures. Features such as document sets, document IDs, and the Content Organizer will help organizations tidy up SharePoint document libraries, while features such as the managed metadata service and folder inheritance can capture document metadata to aid search and navigation. By enabling records management in all SharePoint lists, rather than just in the Records Center, SharePoint Server 2010 could also encourage SharePoint users to play a bigger part in capturing records. The managed metadata service and records management improvements will support centralized SharePoint content management policies and improve consistency.

However, in the end SharePoint Server only provides content management mechanisms, not policies. Organizations will have to define policies for SharePoint content that users will accept and implement those policies with SharePoint to minimize their impact on users while preventing users from circumventing them. Finding the right balance will require understanding of the organization's business and the practices and regulations of its industry, which a generic product like SharePoint can't address. Consequently, organizations that want to benefit from the new capabilities in SharePoint Server 2010 may have to seek help from Microsoft partners that know SharePoint, content management, and their specific industries.

Furthermore, not all organizations will make SharePoint Server their hub for content management and records management. Some have dedicated content management and records management software from other vendors, which they have already adapted to their businesses and policies. Also, some will want to leave at least some records in source applications, such as Exchange, which entails managing them by means other than SharePoint. And some will need content management functions that SharePoint does not provide out of the box, such as imaging and metadata capture for paper documents.

Consequently, there remains room for content management software vendors to work alongside SharePoint. In fact, many vendors offer integration technologies for working with SharePoint sites and content, and most will update these technologies for SharePoint 2010. For example, EMC's My Documentum product enables users to access content in its Documentum content management software via SharePoint sites; its SourceOne product enables organizations to archive and enforce policy for content in SharePoint as well as Exchange and other repositories; and its Captiva software enables capture and validation of physical documents and metadata into SharePoint. All of these products have already been updated for SharePoint 2010 or are slated to receive updates by the end of the year.

Future improvements to SharePoint will include the Content Management Interoperability Services (CMIS) interface, a protocol for communicating with content management systems that Microsoft has been jointly developing with EMC, IBM, and other vendors. If implemented by these vendors, CMIS could ultimately help organizations integrate content from disparate source systems, including SharePoint Server, into overall content management systems.

Web Content Management Improved in SharePoint 2010

SharePoint 2010 updates Microsoft's premier Web content management product (for creating and managing pages and content on a Web site) with greater browser compatibility, accessibility, better use of metadata, and more granular controls for enforcing visual standards while simplifying tasks for content creators. The product will be welcomed by existing customers who want to move past previous constraints on site size and need more powerful content management tools to do so. However, SharePoint's Web management tools are not as mature as its collaboration capabilities, and customers with very large sites will likely find the remaining constraints problematic.

Web Content Management

Even small Web sites face content management challenges, as content producers generate new or refreshed content and try to fit it into a site's navigation menus and visual standards. For large corporate sites, the problems are multiplied, as design is usually delegated to one group while others generate the content. Content creators often need access to corporate logos and other design elements that give the site a consistent look and feel. Organizations also want to maximize reuse by converting print to Web content and vice versa.

Web content management (WCM) systems are designed to manage this complexity, making it easy for content producers to create or update any part of the site for which they have responsibility and giving them access to content and resources (such as images or video) produced by others, while ensuring that the pages they create meet corporate design standards, link to the appropriate items, and are properly indexed for searches.

Most systems let designers and programmers, who are focused on the visual and interactive elements of a site, create templates for Web pages on which some elements are locked to ensure cross-site consistency, while enabling other parts of these pages, such as text or image boxes, to be revised or modified by content creators, such as subject matter experts. Revision histories ensure that changes can be reversed and are visible for other purposes, such as project management. WCM systems also help organizations reuse content by giving them easy ways to convert printed text and graphics for use on the Web or in PowerPoint presentations.

Microsoft's initial entry into WCM was Content Management Server (CMS), based on the Resolution product acquired in the 2001 purchase of NCompass Labs. CMS functionality was rolled into SharePoint Server 2007, eliminating overlaps and allowing the company to leverage SharePoint's much larger partner community and market. The WCM features of SharePoint compete with a large number of commercial and open source WCM systems, including Documentum (owned by EMC), Drupal (open source and used by some government Web sites), Interwoven (Autonomy), Stellent (Oracle), and Vignette (OpenText).

Content Creation Interfaces

SharePoint 2010 better supports multiple browsers and Web standards, gives Web site designers more granular control of master pages, and leverages metadata to reduce site maintenance tasks.

Browser and Standards Support

A common pain for designers—ensuring that Web pages and interactive elements display and work properly on different browsers and are accessible to users with disabilities—is reduced by better browser support in SharePoint and SharePoint Designer.

The text editor for user-generated content, for example, works with the Firefox (3.x and higher) and Safari browsers, as well as Internet Explorer 7 (IE 7) or higher. SharePoint now also supports widely accepted standards, such as XHTML and the Web Content Accessibility Guidelines (WCAG) 2.0 AA standard for accessibility for disabled users.

Ribbon Interface

The most visible change to the user interfaces of SharePoint 2010 sites and SharePoint Designer 2010 (a tool that can be used to customize SharePoint sites without writing code) is the addition of the Ribbon, which was introduced in Office 2007.

Content managers with the appropriate permissions can edit content on a SharePoint Web page more easily, since the Ribbon gives them access to text styles and other features that previously may have required a separate HTML editor. User-generated content can be more expressive because even simple text input boxes, such as for entering comments on a Wiki, can use rich text (with different fonts or bold type, for example).

The Ribbon is also used in SharePoint Designer 2010 and exposes more features to Web site designers. Developers can extend it with new tabs and menu items, including custom actions that initiate a workflow, show a form, or take the user to another Web page.

Designer Controls

SharePoint 2010 gives organizations greater flexibility in delegating content creation and administration. This is important for large SharePoint farms where subject matter experts or business units need to contribute content, but may not have the skills to format this content in HTML or other Web-appropriate formats, leading to inconsistent quality and bottlenecks in keeping the site current.

Specific improvements include the following:

Document sets, a new SharePoint 2010 feature, are folders that define templates for specific content types, such as a Web page, and that can be made specific to a particular group or department. While templates—starting pages that include common page elements, tables, and other design features—are not new, document sets offer greater granularity and allow a department to have a set of Web page templates that incorporate elements that are locked for editing, such as organization-wide logos or colors, along with elements that the department can change, such as headlines, text, and images.

SharePoint Designer 2010 is a more complete SharePoint companion than its predecessor, with greater power to manage SharePoint data structures, and it could reduce the need for a full-blown coding environment, like Visual Studio, for Web page development. (For an illustration, see "SharePoint Designer Content Management Interface".) Unlike its predecessor, Designer 2010 can create, modify, and delete columns, lists, and content types. These are the foundations of SharePoint content management, and they can be accessed by Web site designers to dictate what content appears on a Web page and how it is displayed. Designer 2007 was the preferred tool for creating SharePoint 2007 workflows; Designer 2010 can also import workflows from and export workflows to Visio 2010.

Designer 2010 only works with SharePoint 2010; it can't edit SharePoint 2007 sites.

Digital media controls let Web page designers manage many settings associated with images or video, some of which previously might have required using separate editing software. For example, users can locate and specify digital media, such as pictures or video, and set variables such as the video preview frame, frame size, and whether the video loops or runs one time.

Master pages in SharePoint 2010 can be designed with more granular change permissions. For instance, individual pages may be locked down so that users can only change specified text fields, or can be opened up for users to add new pages and fields and change content types. The designer can specify how a particular content type will be displayed in a field or Web Part, and the libraries from which content for a field can be selected, making it easier for content creators to locate content while ensuring that they use content only from the appropriate library.

Developers and Metadata

SharePoint stores documents in SQL Server by default, although third parties can extend SharePoint to work with stored documents elsewhere, such as in a file system. Most SharePoint content in each site is stored in tables called lists, each row of which is called an item. For example, a SharePoint calendar is a list of event items.

Using a database rather than a file system has the advantage of allowing descriptive document properties (document metadata), such as document title, creator, and the last modified date, to be associated with the document in metadata "columns." This metadata can be used in many ways, such as in search (identifying all documents associated with a particular department or author, for example) and document management (classifying documents according to an organizational taxonomy).

Designers can take advantage of SharePoint's metadata management features when designing Web sites, because metadata can be used to populate columns in SharePoint lists and fields in Web pages. For example, most pictures taken by digital cameras use an exchangeable image file format (EXIF), which SharePoint 2010 can expose for use on SharePoint Web pages, such as displaying thumbnail images or picture captions embedded in EXIF data.

Web Parts can also access metadata. SharePoint's Content Query Web Part can be used to display lists of items based on values of specific columns (e.g., a list of all of the documents associated with a particular author, regardless of content type) as well as metadata associated with the page on which the Web Part is placed. When combined with custom templates, this can become a powerful tool for automatically populating Web page elements.

For example, a Web site designer can specify that every time a content editor uses a template for a new subsite devoted to a city, the editor must specify the city name in a Destination metadata field that will be applied to all content on all pages in the subsite. Thus, if a user at a travel site enters information about a restaurant in Sydney, Australia, in a text box on a wiki page built with a subsite template, that text would be saved with "Sydney" as the value in a Destination column. A Content Query Web Part on a different page in the Sydney subsite could be programmed to automatically display wiki entries labeled with the same Destination column (Sydney) as the page on which the Web Part is used. This gives the site designer an easy way to populate that page with relevant, up-to-date content.

Other Improvements

Other improvements to SharePoint 2010 include better export of content to production sites and better built-in Web analytics.

Content export. Moving content from staging to production has been difficult for large SharePoint sites. Changes made to pages while the site was being exported, such as a page being moved while the export was taking place, could lead to inconsistencies, such as invalid links, or even export failure. When used with SQL Server 2005 Enterprise or later, SharePoint 2010 can use database snapshots, static tables that represent the database at a particular point in time. By specifying the snapshot rather than the active database as the source for the export, the export will be more reliable, and editors or designers can continue to work on the site without waiting for the export to conclude.

Enhanced usage reporting. SharePoint 2010 delivers an expanded system for usage reporting (now called Web Analytics), which enables organizations to track usage of specific pages and features in the product. Among other things, SharePoint 2010 provides more options for filtering reports (for example, by date range). Many reports can be exported to Excel for further analysis. SharePoint's built-in usage reporting still focuses on aggregate usage statistics such as hit counts; to track typical user paths through a site or identify user segments from such patterns, organizations will still need a custom or third-party Web analytics tool.

Web Trails Collaboration Benefits

In spite of the advances in SharePoint 2010, SharePoint continues to be more valuable as a collaboration and enterprise document management tool than as a Web development tool.

While SharePoint is a powerful tool for storing, searching, and accessing Office documents, Office cannot easily be used to create HTML content for a Web site—for example, Word cannot use a template created in SharePoint Designer to create or edit a Web page. Tools for converting Office documents to Web formats such as HTML are limited and focus primarily on archiving documents. As a result, organizations that create most of their documentation with Office must copy and paste from Office to SharePoint or SharePoint Designer (although less reformatting is required in SharePoint 2010 than with its predecessors) to ready it for use as an HTML page on a SharePoint site.

Among other benefits, Web pages are better when an organization needs to ensure "one version of the truth." Any changes to a human resources policy that is accessed via a Web browser, for example, will be immediately available across the organization. In contrast, a similar document that is downloaded to the user's computer will not automatically incorporate later changes, and multiple versions of the document may exist in the organization.

This is not a drawback to organizations that value SharePoint primarily as a document repository and internal search engine, and that rely on users to download or check out documents to local computers. But it means that companies that want to repurpose for their Web site any documentation created by common Office applications may need more advanced Web content management tools.

Improved Dashboards, Reports from SharePoint

Corporate performance management, reporting, data analysis, and other business intelligence (BI) functions are noticeably improved in SharePoint Server 2010. BI technologies aim to help users make better decisions and spot problems early by analyzing data gathered from applications used to manage the business day-to-day. SharePoint could simplify BI with a single authoritative set of tools for users and a single infrastructure for managing those tools and connecting them to data. However, overlapping technologies make it difficult to choose the right platform for long-lived solutions.

SharePoint for Business Intelligence

SharePoint 2010, like earlier versions, includes features for designing and hosting Web-based reports and data analysis tools such as spreadsheet models. Microsoft's key BI product is SQL Server, which can extract, summarize, and analyze large volumes of data from a variety of sources (including non-Microsoft databases) and generate reports. However, Microsoft has been promoting SharePoint as its primary delivery vehicle for the end products of analysis from SQL Server and other sources: a single, authoritative place for users to access reports, spreadsheets, and other analysis tools, and a single infrastructure for IT to manage and update those tools.

Specifically, several popular desktop applications, including Excel 2007 and 2010, Access 2010, and Visio 2010, enable expert users to build interactive reports and charts from SQL Server and other data sources and publish them to SharePoint users who might not have the corresponding applications on their own PCs. SQL Server Reporting Services can also deliver Web-based reports via SharePoint (although Reporting Services can also be used stand-alone).

As a result, organizations that have already deployed SharePoint for team collaboration, corporate portal hosting, enterprise search, or other functions, can extend those deployments for BI, leveraging existing SharePoint hardware, licenses, and expertise. They can use SharePoint security to control who can change data and reporting methods, and to publish selected reports to external users (for example, reporting on inventory levels to suppliers). Microsoft itself is exploiting SharePoint's BI capabilities in a variety of other products, including Project Server, the Dynamics business applications, and the upcoming Duet Enterprise platform for access to data in SAP applications.

SharePoint 2010 does not have a single BI feature, but several overlapping ones. The three most important ones discussed here are the following:

- PerformancePoint Services, which supports business scorecards and other corporate performance management functions
- Excel Services, which delivers Excel spreadsheets to SharePoint users in a browser

- Reporting Services integration components, which enable SQL Server Reporting Services reports to be published in SharePoint.

Several technologies not covered here in detail also aid BI with SharePoint Server. PowerPivot enables assembly and analysis of large data sets in Excel 2010 and SharePoint Server 2010. Visio Services enables publication of Visio 2010 diagrams that incorporate dynamically updated application data, such as a data center diagram that shows system status with data from a systems management application. (See the illustration "Visio Services".) Access Services enables creation of SharePoint-hosted Web applications with Access 2010; it would normally serve departmental applications for tasks such as asset tracking, but it might be used specifically for reporting. (See the illustration "Access Services".)

For Microsoft and its partners, SharePoint's BI features have several attractions. They promote upgrades to versions of desktop applications (such as recent versions of Excel) that provide extra features for SharePoint BI and encourage cross-sales of SharePoint to organizations that own those desktop applications. They expand the audience (and license base) of SQL Server and other applications used as data sources by bringing in SharePoint users. They also increase revenue from SharePoint itself: except where indicated below, SharePoint 2010 BI features require the full SharePoint Server product (not the free SharePoint Foundation) and its Enterprise Client Access License (CAL) for every client.

Performance Management Merged In

SharePoint Server 2010 incorporates PerformancePoint Services, a set of design tools and components to help organizations track business performance on an ongoing basis and flag potential problems (a task sometimes called *corporate performance management*). PerformancePoint Services delivers Web-based reports called *dashboards* that provide performance summaries for an organization, process, job role, or project. A dashboard usually includes a *scorecard*, which delivers a collection of metrics called key performance indicators (KPIs) to provide a complete, up-to-the-minute view of performance. However dashboards can also include Web-based analysis tools and reports to shed additional light on the scorecard. Reports can come from a variety of sources, including SQL Server Reporting Services and Excel spreadsheets. Analysts design SharePoint dashboards in the PerformancePoint Dashboard Designer client, a specialized tool for defining KPIs and structuring scorecards and dashboards. (See the illustration "PerformancePoint Dashboard Designer".)

Data for PerformancePoint dashboards can come from many different sources, including SQL Server, SharePoint itself, and other business applications via SharePoint Business Connectivity Services, an interface that enables users to view, update, and search selected data in sources external to SharePoint, such as customer relationship management systems. For dashboards to be useful, organizations generally need to design data sources with corporate performance management in mind and update them regularly. An ideal data source is a corporate data warehouse or business unit "data mart" that incorporates cleaned-up, summarized snapshots of application databases on a regular basis. Setting up and maintaining such data sources is often more challenging than designing and delivering the dashboards themselves. However, if an organization is already maintaining useful data sources, PerformancePoint dashboards could help deliver relevant data from them to users.

PerformancePoint Services is based largely on the Monitoring and Analytics components of PerformancePoint Server 2007, a discontinued product. PerformancePoint Services delivers several improvements over its predecessor.

Integration with other SharePoint Server features and infrastructure. PerformancePoint Services is a SharePoint Server 2010 service application, which means it can be deployed and administered with SharePoint Server. Furthermore, all PerformancePoint KPIs, scorecards, dashboards, and other data are stored in SharePoint Server content databases, and access to them is controlled with SharePoint security. PerformancePoint content can also be incorporated into other types of SharePoint pages with other Web Parts, such as the new Chart Web Part. (See the illustration "SharePoint Chart Web Part".) Consequently, PerformancePoint Services dashboards should be easier to deploy and maintain than dashboards created with PerformancePoint Server 2007. With PerformancePoint Server 2007, dashboards and other content were deployed to users via SharePoint sites, but organizations had to maintain a separate PerformancePoint Server database and infrastructure to store master copies for dashboard authors.

Dashboard, scorecard, and KPI enhancements. PerformancePoint Services delivers many new design options for dashboards and their components that will make dashboards easier to set up for analysts and help end users understand what they are seeing. For example, dashboards can incorporate pie and decomposition tree charts to help analyze breakdowns of particular KPIs. (These "analytic charts" work only with data in SQL Server Analysis Services databases.) Scorecards can include hierarchies that break down KPIs into components and that are updated dynamically. For example, a scorecard can include a breakdown of a KPI by a geographic region hierarchy that is imported and updated from the data source. PerformancePoint Server 2007 scorecards could be based on hierarchies, but analysts had to define the hierarchies and maintain them manually; they were not derived from data sources. Also, PerformancePoint Services can calculate KPIs from an analyst-supplied formula and data drawn from more than one data source; in PerformancePoint Server 2007, each KPI had to draw on a single data source and calculations had to be performed by the source database.

Migration Tricky, Overlap with Built-In KPIs

Despite these and many other improvements, PerformancePoint Services is not compatible with all dashboard content from PerformancePoint Server 2007; for example, PerformancePoint Services cannot import PerformancePoint Server 2007 PivotTable or PivotChart views. This means analysts might have to manually re-create features of existing PerformancePoint Server 2007 dashboards to migrate them into PerformancePoint Services. Also, PerformancePoint Services does not completely replace PerformancePoint Server 2007. In particular, it does not include PerformancePoint Server 2007 Planning, a financial planning and budgeting application.

SharePoint Server 2010 also has a completely separate feature for defining KPIs and scorecards, called *status indicator lists*, which is based on the KPI lists feature of SharePoint Server 2007. This feature has a minimal Web-based forms interface for defining KPIs (now called status indicators) and connecting them to a variety of data sources. The feature has changed little from the previous version and offers only a subset of the capabilities of PerformancePoint Services for scorecards.

Excel Services Keeps Up with Desktop Excel

SharePoint Server 2010 delivers an updated version of Excel Services, the product's component for sharing Excel workbooks with users who might not have or use Excel. Excel Services runs workbooks and performs calculations on SharePoint Server, displaying results to users through a browser. Users can make temporary changes to workbooks such as filtering and sorting tables and entering selected input values (called parameters) to do what-if analysis, but they cannot permanently change the source workbook. Excel Services is designed for interactive models that are regularly used by large audiences, such as a pricing tool for salespeople or a proprietary model for evaluating investments. For these types of workbooks, Excel Services has advantages over sharing via file servers or e-mail: it centralizes formulas and data, which makes it easier to keep the spreadsheet current and prevent unauthorized changes, and it does not require Excel on users' PCs. However, Excel Services offers only a subset of Excel features. It cannot open some Excel workbooks, and it can't display all features of workbooks it can open.

Some notable changes to Excel Services in SharePoint 2010 include the following:

Editing and collaboration through a separate Excel Web App. SharePoint Foundation 2010 and SharePoint Server 2010 can host Excel Web App, an add-on feature that is separate from Excel Services but shares some components. Like Excel Services, Excel Web App enables browser users to view Excel workbooks stored on a SharePoint Server, but Excel Web App also enables users to make permanent changes to formulas, data, and formats. Excel Web App also enables several users to edit a workbook simultaneously and see one another's changes. Excel Services, in contrast, does not enable users to make permanent changes, except by editing a workbook in the Excel client and then republishing the result to SharePoint, and in Excel Services each user gets his own, isolated view of a workbook. The two features also are licensed differently: Excel Services requires SharePoint Server as well as an Enterprise CAL for each client, while Excel Web Apps requires SharePoint Foundation or SharePoint Server as well as an Office 2010 Standard or Professional Plus license for each client.

PowerPivot data analysis. A SharePoint Server 2010 add-on enables Excel Services to process Excel 2010 workbooks that contain PowerPivot databases. PowerPivot enables Excel 2010 users to create and process large multidimensional databases (sometimes called cubes) embedded in Excel files and build Excel tables and charts for analyzing those databases. (See the illustration "PowerPivot Components".) The technology enables authors who are very familiar with Excel and the data to create sophisticated analysis solutions with much larger data sets than Excel and Excel Services could normally handle, with minimal IT or developer support. (The SharePoint PowerPivot add-on is delivered as a special "SharePoint integrated" installation mode of SQL Server 2008 R2 Analysis Services.)

Excel 2010 features and better compatibility. In SharePoint 2010, Excel Services has been updated to work with several new Excel 2010 features, including Sparklines (small line charts that can be embedded in a cell) and the Slicers filtering user interface. Excel Services also offers better compatibility with Excel generally. For example, it can open and display Excel workbooks that contain embedded images (which are common in reports designed with Excel), and it can open workbooks that

contain Visual Basic for Applications (VBA) code, even though it will not execute the code. As before, users will receive a warning when Excel Services opens a workbook that contains incompatible features (such as VBA code).

New API for JavaScript and Silverlight clients. Excel Services in SharePoint 2010 offers a new REST Web services API. REST (Representational State Transfer) is a style of Web services API widely used on the public Internet that uses the standard HTTP protocol and is designed to be readily accessible to applications written in client-side Web technologies such as JavaScript or Silverlight. The new API could be particularly useful for creating Web pages that leverage an existing Excel workbook on SharePoint Server; for example, an organization might make parts of a shipping costs calculator available to its customers online.

Excel Services spreadsheets and charts can be incorporated into PerformancePoint dashboards, and PerformancePoint KPIs can also be calculated from data in Excel spreadsheets on SharePoint. Among other things, this could help organizations that currently use Excel for performance management to move to PerformancePoint Services. However, Excel features such as conditional formatting and embedded charts can also be used to build dashboards directly, without PerformancePoint Services, although its specialized Dashboard Designer might make the task easier.

Reporting Services Updated to 2008 R2

SQL Server Reporting Services integration components have been updated for SharePoint Foundation 2010 and SharePoint Server 2010. These components enable users to publish Reporting Services Web-based reports to SharePoint for distribution to other users. Among other things, this enables organizations to control access to reports with SharePoint security and to exploit an existing SharePoint installation for reporting as well as other functions, such as team collaboration and search. (As noted earlier, Reporting Services integration can work with SharePoint Foundation; it does not require the full SharePoint Server or its Enterprise CAL.)

SharePoint 2010 gets a few improvements for Reporting Services integration. It is compatible with SQL Server 2008 R2 Reporting Services, the latest version. Among other things, that means SharePoint 2010 can exploit 2008 R2 features such as the ability to generate reports from lists stored in SharePoint.

However, SharePoint 2010 is not compatible with SQL Server 2005 Reporting Services. Specifically, it requires SQL Server 2008 or later for the report servers that host and render reports, and SQL Server 2008 R2 for the integration components that run on SharePoint Web servers. Among other things, this means that organizations moving to SharePoint 2010 will have to upgrade servers and reports to the appropriate versions of SQL Server Reporting Services. Most existing reports will upgrade automatically, but some might require changes.

Reporting Services reports, like Excel Services workbooks, can be embedded in PerformancePoint dashboards. However, Reporting Services can also be used on its own to create dashboards, and SQL Server 2008 R2 includes some new features (such as graphic icons for KPI status) for scorecards as well.

Caution Warranted for Long-Lived Solutions

SharePoint 2010 delivers valuable improvements for BI functions, but it also increases the number of overlapping technologies for specific tasks. For example, to create a dashboard with a scorecard and some embedded charts for SharePoint Server 2010 users, a developer might use one or more of the following technologies:

- PerformancePoint Services
- Status indicator lists and the Chart Web Part
- Excel Services, with or without PowerPivot
- Reporting Services

Each of these technologies has its own advantages and drawbacks. Organizations trying to build a short-lived solution as quickly as possible can choose among these technologies strictly on merits and on what expertise they have available; for example, users familiar with Excel will naturally gravitate to Excel Services, while professional developers already working with the SQL Server BI platform will naturally choose Reporting Services.

However, organizations building long-lived solutions should consider how much development effort Microsoft will give each technology in the future. If history is any guide, not all will get equal effort, and some might eventually be dropped, as was PerformancePoint Planning. In the long term, solutions on discontinued or de-emphasized technologies become increasingly difficult to maintain as bugs go unfixed, compatibility problems appear with related technologies (such as new versions of SQL Server), and the supply of trained developers and administrators dries up.

Several factors help predict whether a technology will continue to attract Microsoft investment and survive in the long term: which Microsoft organizations are responsible for it, how much development effort it has attracted historically, how technically stable it has proven from release to release, and how many customers (including Microsoft business units) have adopted it. On these criteria, Reporting Services and its SharePoint integration components seem most secure: Though relatively new, they are backed by two powerful Microsoft organizations (the SQL Server and SharePoint units), they have received steady effort and maintained a fair record of technical stability, and they have many Microsoft and external customers already. PerformancePoint Services, in contrast, is new to the SharePoint organization and descended from a series of discontinued products (Business Scorecard Manager 2005, PerformancePoint Server 2007), although it has received steady improvement over several years. It has fewer existing customers than Reporting Services both inside and outside Microsoft. In the middle lies Excel Services, which has powerful backers (including the Office and SharePoint units) and some internal customers (such as Microsoft's Project team). It is still relatively new, but its link to desktop Excel gives it a large base of potential users.

New SharePoint Aims for Larger Scale

SharePoint 2010 delivers simpler administration, reduced downtime, and improved scalability. As a result, organizations should be able to scale SharePoint 2010 installations to larger numbers of users and volumes of content with less effort than was required in earlier versions. However, most organizations building large-scale SharePoint installations will still need third-party administration software to keep their installations in working order, and for large-scale document sharing and content distribution SharePoint still lacks some capabilities of file shares and Exchange public folders.

Administration: Improved Design Options, Automation, Monitoring

SharePoint 2010 includes many improvements to its infrastructure and management tools to ease administration. The improvements will particularly benefit large-scale organizations whose SharePoint installations serve many internal stakeholders.

Service Infrastructure Expands Choices for Architects

SharePoint 2010 has an improved infrastructure for managing installation of its most important supporting services, giving organizations more fine-grained control for scaling out SharePoint installations and for striking a balance between security and simplicity.

SharePoint installations depend on a variety of supporting services, such as the following:

- A user profile service, which manages SharePoint user data to enable content personalization, personal Web sites, and social networking features
- Search query and indexing services, which enable keyword search of SharePoint content and external data sources, such as file shares
- The Excel calculation service, for viewing and recalculating Excel spreadsheets in a browser
- The secure store service, which maintains credentials for access to resources outside of SharePoint, such as application databases being accessed via SharePoint sites.

In SharePoint 2010, each major supporting service is packaged into a component called a service application, which an organization can install on a server or group of servers, separately from the Web servers that run SharePoint Web applications. A service application can be shared by multiple SharePoint farms, multiple Web applications within a single farm, or be restricted to a single Web application in a farm. (For a review of architecture terms and an example, see the illustration "Anatomy of a SharePoint 2010 Farm".)

Service applications resemble the Shared Service Provider (SSP), a component that delivered supporting services in earlier versions of SharePoint. Like an SSP, service applications can be run and scaled separately from Web servers. Service applications also isolate services for security purposes: for example, a business unit with sensitive data (such as a human resources department) might deploy its own service applications for searching those data. Conversely, service applications

enable centralized management of services that work with shared data; for example, an organization might use a single, global user profile service application for all its SharePoint farms.

The main difference: an SSP included all SharePoint supporting services, not just a single service, and those services had to be taken as a group. This complicated management and scaling. For example, the only way to give a department its own isolated search service was to give it its own SSP, complete with user profile and other services, which would then have to be synchronized with an organization's primary SSP.

SharePoint 2010 service applications have other benefits over SSPs, including the following:

- New APIs enable developers to create custom service applications that extend the capabilities of SharePoint; for example, Microsoft's own Project group has used this capability to integrate its project management software into SharePoint Server
- Service applications can be administered with SharePoint 2010's Central Administration Web interface; SSPs used their own administrative console, requiring administrators to learn and switch between two consoles
- The free SharePoint Foundation includes the service application infrastructure, which means administrators and developers can use the same interfaces as in the full SharePoint Server; however, SharePoint Foundation has a much smaller set of built-in service applications than SharePoint Server.

Service applications do complicate the job of SharePoint architects. SharePoint 2010 has many more services than previous versions, both because it has more features (such as the Office Web Apps for Web-based editing of Office documents) and because services from earlier versions have been divided into multiple service applications. As a result, it is less obvious with SharePoint 2010 which service applications should be installed, and where. Microsoft has already begun to publish architecture guidance for the product, but architects will have a lot to learn.

Multitenancy Supports Hosting, Decentralized Management

SharePoint Server 2010 has built-in support for multitenant installations: a single farm can host multiple independent organizations (tenants), each with its own administrators who use a tenant-specific console to manage all SharePoint content, service application data, user accounts, and other resources. A separate group of farm administrators retains control over the tenants and farm administration tasks, such as patching SharePoint and backing up SQL Server databases.

The built-in multitenant capability will simplify setup and management for firms that offer multitenant SharePoint hosting, including Microsoft with its SharePoint Online Standard service. The capability could also help large IT organizations, enabling them to delegate most SharePoint Server administration to business unit IT personnel as tenants, while retaining control of the shared farm infrastructure.

However, not all SharePoint Server 2010 service applications and features are multitenant. The most notable omissions are FAST Search for SharePoint, a set of enhanced enterprise search components, and the PerformancePoint Services business scorecarding feature. While hosters might be able to create multitenant services that include these features (Microsoft intends to), they will have to develop their own infrastructure to do so.

For SharePoint service applications that are multitenant, there are also limitations. Some SharePoint features are not tenant-aware; for example, backup utilities do not provide a way to back up or restore data for specific tenants out of the box. There is also no graphical user interface for setting up multitenant installations initially. Developers can use PowerShell scripting and the SharePoint API to close these gaps.

PowerShell Leads Other Administrative Improvements

Several other features will help administrators maintain SharePoint, including the following:

PowerShell 2.0 and expanded command set. SharePoint 2010 adopts PowerShell 2.0 and adds substantially to the product's PowerShell scripting command set, enabling administrators to perform virtually any operation with the latest version of Microsoft's administrative scripting technology. PowerShell scripting gives an organization a way to automate and standardize its SharePoint administrative procedures and communicate them to its administrators. Furthermore, as noted above, some tasks (such as initial setup of multitenant farms) can only be done in PowerShell. Earlier versions of SharePoint enabled some operations with PowerShell, but some tasks required the SharePoint command-line administration tool (stsadm). That tool remains available, but it is deprecated; Microsoft will not update it for future SharePoint capabilities and might remove it in the next version.

However, SharePoint's console still lacks some PowerShell features of other Microsoft products, such as Exchange Server. For example, the SharePoint Central Administration console cannot generate and save PowerShell scripts for administrative operations done at the console, a useful way to capture an instance of a procedure for repeatable use in a PowerShell script.

Managed accounts. An organization can opt to let SharePoint 2010 itself manage the passwords of SharePoint service accounts. For such managed accounts, SharePoint can assign a random password and change it automatically on an administrator-set schedule, or notify an administrator to change it manually. The passwords are kept in an encrypted store and can be reset only by a SharePoint farm administrator. Administrators can also see which SharePoint services are running under a given managed account. The managed accounts feature will help organizations maintain more "least privileged" configurations, in which each service runs in a security account with only the privileges it requires. Such configurations are more secure, but also lead to proliferation of security accounts and passwords; the managed accounts feature helps administrators deal with this proliferation.

Note that Windows Server 2008 R2 has a similar feature for Active Directory service accounts, but the SharePoint 2010 feature is separate and does not require Windows Server 2008 R2.

Improvements for Health and Usage Tracking

SharePoint 2010 delivers several improvements to help administrators monitor the health of their installations and learn which features and content are most used.

Web Analytics enhanced usage reporting. SharePoint 2010 delivers an expanded system for usage reporting (called Web Analytics in SharePoint 2010), which enables organizations to track usage of specific pages and features in the product. Among other things, SharePoint 2010 provides more options for filtering reports (for example, by date range). Many reports can be exported to Excel for further analysis. SharePoint's built-in usage reporting still focuses on aggregate usage statistics such as hit counts; to track typical user paths through a site or identify user segments from such patterns, organizations will still need a custom or third-party Web analytics tool.

Health Analyzer. A tool similar to the SharePoint Best Practices Analyzer, which spots common misconfigurations and other problems and recommends changes, has been integrated into the SharePoint 2010 administrative user interface. Called Health Analyzer, this feature is useful both for new installations and to catch later configuration changes that violate an organization's policies. The Health Analyzer uses a set of policy rules that organizations can enable or disable, and developers can create new rules with the SharePoint 2010 API.

Developer Dashboard. Organizations can enable a control called the Developer Dashboard on a Web page to show performance statistics and runtime error information. The control can be turned on globally for a farm or page-by-page by users with appropriate permissions. The control helps developers debug and optimize pages they have customized, but it could also help administrators diagnose pages that load slowly or generate errors. Enabling the control is simpler than instrumenting a page with tracing or attaching a debugger, although these options are still available to developers who want to probe further.

Logging Infrastructure Enhanced

SharePoint 2010 delivers improvements to its unified log database and infrastructure, which captures log entries for SharePoint and underlying components. Notably, the log database schema is documented and Microsoft will support access by custom and third-party applications. This enables organizations and partners to design their own reports and analysis tools to supplement SharePoint's own. Specific improvements include the following:

- SharePoint 2010 assigns a unique event correlation ID to each Web request it receives and includes those IDs in logs and error messages returned to users; this helps administrators correlate user error messages to logs and diagnose problems
- Logs capture usage of specific SharePoint features, such as Excel Services, although SharePoint's built-in usage reports do not display this information; among other things, feature usage can help organizations determine whether they need SharePoint's enterprise Client Access Licenses (CALs) for their users

- By default, SharePoint 2010 will throttle event logging to limit the rate at which events are captured, which will help prevent overload of servers during bursts of logging activity (for example, after a serious error has occurred).

Note that the log database described here captures overall health and usage of SharePoint application components and accesses to pages, but it does not record every access to every stored document by each individual user, so it cannot be used to monitor compliance with policies or regulations that govern information access. SharePoint 2010, like earlier versions, can maintain item-level audit logs that do record each access operation; those logs are stored in content databases and managed with a different set of tools than the health and usage logs.

Recovery and Patching to Reduce Downtime

Several SharePoint 2010 features reduce downtime by speeding recovery after failures, simplify backup and recovery of data, and reduce the impact of maintenance operations. These features reduce the effort organizations need to put into high-availability SharePoint installations, even if they do not fundamentally change the product's capabilities.

More Automated Mirroring

SharePoint 2010 offers better tools for maintaining "warm spare" secondary database servers with SQL Server database mirroring, and it can automatically fail over to a secondary server when the primary server becomes unavailable. SQL Server database mirroring uses one-way replication of databases from a primary server to read-only copies on a secondary server. Mirroring in SharePoint helps organizations provide temporary, read-only service during maintenance or after loss of a primary server or its data center. Previous versions of SharePoint could use SQL Server mirroring, but SharePoint provided no setup tools for mirrored configurations and administrators had to script their own failover procedures. Note that SharePoint 2010's mirroring feature requires SQL Server 2008 (or later) configured for automatic failover and its synchronous "high-safety" mirroring mode, which ensures that data are copied to the secondary server on each transaction so that it never lags behind the primary server. Automatic failover and high-safety mirroring entail extra hardware, including a witness server to control failover, and a fast network connection between the primary and secondary database servers.

SharePoint 2010 also will work more gracefully with read-only content databases after a failover than previous versions did; for example, when it detects that a site's content database is read-only, it will automatically disable user interface commands that would enable users to write content.

In SharePoint 2010, as in earlier versions, organizations can also protect database servers with SQL Server failover clustering, which provides automatic failover between servers that share disk storage. Failover clustering provides faster recovery than disk mirroring after loss of a server and is mostly transparent to SharePoint and its users; mirroring is slower and provides only read-only access after a failure. However, failover clustering requires even more specialized (and expensive) hardware than high-safety mirroring with automatic failover, and it maintains databases on storage shared by the primary and secondary servers,

which means it can't protect against failure of the storage system itself. Failover clustering is also more difficult to implement for disaster recovery, where servers are kept apart from one another in distant data centers.

While mirroring enables an organization to maintain secondary database servers and farms, it is not designed for replicating content across farms. In particular, it's not a way for a global organization to maintain a single, shared collection of content (such as its intranet portal content) across a set of globally distributed SharePoint farms. SharePoint 2010 retains the content distribution tools introduced in earlier versions, which can help organizations move content from staging or test farms to production ones, but cross-farm replication in general requires additional software, already available from a variety of vendors.

Restores Simplified for Both Content and Farms

SharePoint 2010 delivers several improvements in its built-in backup and recovery feature. Notably, it enables an administrator to restore a SharePoint content database to an unattached database server that is not part of any SharePoint farm and extract sites, lists, or other content from it. This enables restores of individual sites or lists without performing a full restore and is particularly valuable for recovering an accidentally deleted or corrupt site or list that is not available in SharePoint's Recycle Bins. With previous SharePoint versions, administrators could restore sites or lists using Microsoft's Data Protection Manager backup product, but it required a complete, single-server SharePoint installation (called a recovery farm) to restore to. Some third-party backup products also enabled restores without a recovery farm, a capability that SharePoint 2010 now delivers out of the box. SharePoint 2010's built-in backup still provides no administrative or end-user interface to restore an individual document from a content database—it can restore only entire sites and lists. The product's APIs do enable single-item restores, and single-item restore is available in Data Protection Manager 2010 and other third-party backup products that support SharePoint 2010.

SharePoint 2010 backup delivers other improvements. Notably, it can restore an entire backed-up farm configuration to new hardware, including Web application and service application settings. This will be valuable not only for disaster recovery but also for creating copies of farms for staging or testing. Earlier versions of SharePoint could back up farm configurations but lacked the ability to restore them to other hardware. Microsoft also says it has improved the performance of backup for the search service and reduced its impact on ongoing operations.

Nevertheless, organizations managing large installations will probably have to rely on Data Protection Manager 2010 or third-party SharePoint backup products. These products enable incremental backup, which captures only changes since the last backup to reduce space requirements and limit impact on running systems, as well as single-item restores. They provide better tools for cataloging and locating backups, and also plug holes in SharePoint 2010 farm configuration backup, which still does not capture all aspects of a farm; for example, it does not pick up configuration files on Web servers that have been edited outside of SharePoint.

Deferring Downtime for Patches

With SharePoint 2010, Microsoft intends to improve the way that patches update SharePoint databases, enabling organizations to get critical patches applied more quickly while reducing, or at least deferring, some of the downtime.

SharePoint patches frequently change the product's database schema as well as its code. These changes can cause considerable downtime in installations with large databases, as SharePoint must migrate all data into the new schema. Future SharePoint 2010 patches will offer a limited form of backward compatibility that enables an organization to apply a patch to SharePoint's code while leaving the database schema temporarily unchanged. This will enable organizations to apply a patch quickly (for example, to close a critical security vulnerability), while deferring the patch's database schema changes to a scheduled maintenance period (for example, on the weekend). Note that this promised capability depends entirely on Microsoft's internal development discipline; it will work only if the Microsoft SharePoint team consistently delivers backward-compatible patches.

Improved Scalability, but Limits Remain

A number of improvements increase the number of users and volume of content a single SharePoint 2010 farm can handle. These will enable larger SharePoint installations and could help organizations consolidate SharePoint farms to simplify management.

Increased capacity limits and recommendations. Higher hard limits and recommendations for storage size and other thresholds will enable organizations to support larger installations and more content in a single farm. Notably, Microsoft says that SharePoint 2010 can routinely allow users to retrieve up to 5,000 items of a document library or other list with acceptable performance, up from 2,000 in the previous version; this will particularly benefit organizations maintaining large document archives. SharePoint's built-in enterprise search service has doubled the number of items it can index (from 100 million to 200 million), and the new FAST search add-on can handle more than 500 million items. For backup and disk performance, Microsoft recommends that typical content databases be limited to 200GB, up from 100GB in the previous version. As in earlier versions, a SharePoint 2010 farm can store more than 200GB of content by having multiple content databases, but a top-level SharePoint site (called a site collection) must fit into a single content database, so increasing the recommended content database size also raises the maximum recommended size of site collections.

Some fundamental limits and recommendations remain unchanged; most notably, SharePoint 2010 still cannot store individual files larger than 2GB.

Remote binary large object (**BLOB) Storage API**. SharePoint 2010 includes an API for storing files and other BLOB content outside of SQL Server—for example, in a file system. The API requires SQL Server 2008 or later. Storing SharePoint BLOBs outside of SQL Server helps handle large volumes of content more cost-effectively; file system or similar storage is often cheaper per byte than SQL Server storage when all hardware costs are taken into account, and, in some cases, moving BLOBs out of SQL Server will improve performance.

To use the API, organizations need additional software (called a provider) that handles reading and writing BLOBs in the external storage and coordinates backup and restore between the external storage and SQL Server. Microsoft delivers a free provider (called the filestream provider) for SQL Server 2008 R2 that stores BLOBs in the local file system on a database server. However, many organizations will want to store BLOBs on remote file shares or network-attached storage (NAS), which requires a provider from a third party such as Metalogix or NetApp.

Earlier versions of SharePoint included a similar API for third-party storage providers called External BLOB Storage (EBS). EBS has been deprecated, and future versions of SharePoint might not include it. The Remote BLOB Storage API has a number of advantages over EBS; for example, it enables enforcement of SharePoint content retention policies on externally stored BLOBs and provides a PowerShell interface for management operations (such as migrating BLOBs to and from SQL Server). Nevertheless, EBS still works in SharePoint 2010, and some existing solutions built on EBS are being updated for SharePoint 2010. For example, by the end of 2010 EMC plans an update to Documentum Repository Services for SharePoint, its EBS-based product that enables SharePoint BLOB storage in Documentum repositories.

Throttling. SharePoint 2010 can block costly operations that might slow down a system for all its users or render it unstable under high load. For example, access to long lists was a frequent cause of performance slowdowns in earlier versions. By default, SharePoint 2010 will automatically block attempts by users to view more than the recommended 5,000 entries in document libraries or other lists and will return a message telling the user to filter the list. (Administrators can bypass the mechanism for their own access.) SharePoint will also selectively discard new Web requests when resources such as processors, memory, and ASP.NET request queues show signs of overload, using an administrator-defined set of load thresholds. SharePoint drops requests based on a priority scheme that favors writes over reads (to reduce the chance that data will be lost) and interactive users over search indexing. Developers can adjust the priority scheme using the SharePoint 2010 API and PowerShell, but there is no administrative user interface for managing it.

Microsoft says that it has taken several steps to improve SharePoint 2010's responsiveness to users. For example, many SharePoint Web pages have been redesigned to reduce the number of times they hit servers to retrieve page elements, such as images, and execution of some JavaScript code in pages is deferred to enable the user to see the page more quickly. Office 2010 includes a new file transfer client and protocol for SharePoint 2010 that hide delays from users by processing data in the background and transmitting only changed parts of files to SharePoint on save.

Despite scalability and performance improvements, SharePoint 2010 requires more powerful hardware than the previous version. For example, Microsoft recommends 8GB memory and four processor cores for Web and application servers, and 16GB memory for database servers, double what it recommended for the previous versions. However, actual requirements will vary substantially depending on how a server is used (for example, how frequently it writes to databases), and whether it runs demanding service applications (such as the new Office Web Apps) that require more memory and processor resources. SharePoint 2010 administrative documentation provides capacity planning guidelines, including details of some of Microsoft's own internal installations.

Large-Scale SharePoint: Some Assembly Required

Overall, SharePoint 2010 will help organizations deploy and maintain large-scale SharePoint infrastructure centrally, while giving them more options to delegate site and content administration to business units. The version gives architects many more options for scaling out installations and isolating sensitive data and has enhanced built-in database management tools to reduce downtime and speed recovery. The product also raises limits on the amount of content that a single farm can handle, which will help organizations build larger farms or consolidate existing ones to simplify management.

However, large SharePoint installations will continue to rely heavily on add-on software. Even basic operations such as backup will benefit from Microsoft or third-party software, and organizations will need to build or buy software to store SharePoint documents outside of SQL Server, replicate content across farms, and analyze patterns of Web usage in commercial sites or applications, among other tasks. Upgrades from anything other than the most recent SharePoint version will also benefit from third-party software. (See the illustration "Upgrading to SharePoint 2010".) Most of these gaps will not be addressed until the next version of SharePoint, which will probably not arrive for another three years. That leaves a window of opportunity for software vendors who fill the gaps, and for systems integrators to assemble software from those vendors into complete SharePoint solutions.

SharePoint 2010's administration, availability, and scalability improvements also do not fundamentally change the product's most widely used capability, file sharing. In particular, older file-sharing mechanisms such as the Windows Distributed File System (DFS) and Exchange Server Public Folders still enjoy some advantages over SharePoint, such as larger maximum file sizes, replication mechanisms to maintain geographically distributed content, and lower-cost storage. SharePoint will continue to augment rather than replace these mechanisms for file sharing until these gaps are closed.

SharePoint 2010 Improves Tools, APIs

Like earlier versions, SharePoint 2010 enables extensive customization by developers. Customization can range from altering SharePoint's built-in workflows and Web Parts to building entire custom applications on SharePoint's Web platform using the product's APIs. In the middle of the spectrum, developers can integrate SharePoint with external data from applications such as an organization's enterprise resource planning (ERP) system, enabling SharePoint to serve as a single, common access point for the organization's data. SharePoint 2010 delivers enhancements for all of these tasks, including a significantly better access to external data, substantially improved tools for development in the popular Visual Studio development environment, and numerous improvements to the product's APIs and application platform to simplify development and reduce downtime from errors in custom code.

SharePoint 2010 Works Better with External Data

SharePoint and Office 2010 deliver Business Connectivity Services (BCS), a new interface that enables users to view, update, and search selected data in sources external to SharePoint, such as customer relationship management (CRM) and ERP systems. This can make important business systems accessible to more workers on a "self-service" basis, reducing their reliance on back-office specialists to complete day-to-day tasks. BCS replaces the SharePoint Server 2007 Business Data Catalog (BDC) and delivers substantial improvements, such as the capability to write to external data sources and better development tools. However, Office 2010 is required to access the capabilities from within Office applications.

External Data Not New to SharePoint

Previous releases of SharePoint Server have offered connections to external data sources, but these solutions were difficult to implement and offered limited feature sets. Prior to SharePoint Server 2007, external data could only be accessed by custom code written by skilled developers. SharePoint Server 2007 introduced the BDC, which acted as a bridge between the Web portal and back-end databases and business applications, such as CRM or ERP systems.

The BDC stored an XML description (metadata) that instructed SharePoint on how to retrieve information from back-end systems. The data could be displayed in a Web browser through a number of prebuilt Web Parts (user-configurable components that allow Web pages to be easily customized), as well as through SharePoint's enterprise search capabilities. Once the BDC was configured to communicate with a data source, data from that source could appear in multiple places via the BDC Web Parts with no further coding. The BDC also provided a consistent API that custom or commercial Web Parts could use to access data from any supported database or application.

The BDC connected to databases via ADO.NET (the data access library included with the .NET Framework) and to applications that exposed their functionality via Web services. In either case, a developer created metadata that defined types of information, or *entities*, to be retrieved from a database or Web service. A sales tracking database, for example, might contain information on products, customers, and sales orders; the developer would define BDC metadata for retrieving each of those three entity types from the database. The BDC could also make information in back-end systems visible to SharePoint's enterprise search capability if a developer created a special form of metadata giving SharePoint the ability to enumerate all entities in the application (e.g., crawling all customers in a sales database). The BDC neither stored any application data nor processed any queries itself. Instead, it delegated the processing of the request to the back-end database or application.

The BDC had many shortcomings. Developers could specify how information in the back-end system was to be updated, but only in a primitive fashion: a URL would take a user to a site where an entity could be updated. Most database sources don't provide a built-in Web interface for modifying database records, and the BDC did not provide a way to update them directly; thus, in most cases BDC was a read-only interface to external data. To access applications via Web services, the BDC used the generic SOAP Web services protocol; it could not take advantage of the reliability and security extensions to SOAP implemented by the WS-* protocols and Microsoft's Windows Communication Foundation (WCF) API. Also, developer documentation for BDC has been sparse and incomplete, development using the technology required tedious coding of XML metadata, and no

debugging facility was offered. An update to the BDC SDK, subsequent to the release of SharePoint 2007, provided an improved developer tool and some additional documentation, but still presented a difficult developer experience.

SharePoint 2010 Takes External Data More Seriously

SharePoint 2010 replaces the BDC with Business Connectivity Services (BCS), a new set of components that are delivered in SharePoint Server 2010, SharePoint Foundation 2010 (a free edition of SharePoint, formerly known as Windows SharePoint Services), and Office 2010. BCS uses an architectural model similar to BDC but expands its capabilities while dramatically improving its developer experience. Backward compatibility for the BDC in SharePoint 2010 is enabled by a migration procedure and a new Application Registry Service, although some code refactoring may be required.

BCS provides external data access through several connector mechanisms, including the following:

Databases. An ADO.NET connector allows read-write access to any data source with an ADO.NET provider, such as SQL Server.

Web services. A WCF connector allows access to data sources exposed as Web services and WCF endpoints. Access to Oracle, SAP, and Siebel data sources can be accomplished through WCF adapters that are available separately in the BizTalk Adapter Pack.

.NET Connectivity assemblies. Developers can use Visual Studio (VS) 2010 to create a SharePoint .NET Connectivity assembly to enable external data access through SharePoint. A .NET Connectivity assembly is used for access to simple, mostly static data sources for which access through an ADO.NET provider or Web service is unavailable. For example, an infrequently updated product catalog stored in a proprietary format might be accessed using a custom .NET Connectivity assembly.

Custom connectors. Larger-scale data sources with frequently changing data structures are best accessed by developing a custom connector. Custom connectors allow the most flexibility for external data access from SharePoint, but also require the most developer work, including implementation of the data connection, transfer method, and installation of the connector. If an external data source is controlled by an outside party that does not guarantee the stability of its data structures, a custom connector should be used.

BCS continues the BDC's notion of entities to represent data elements. Entities are now called External Content Types (ECTs), and their definitions are stored in the BCS metadata store. An ECT describes how to connect to a data source; perform create, read, update, delete, and query actions on it; and display it in a client. For example, an ECT could define how to access a products table stored in SQL Server, allowing it to be exposed as a list in SharePoint and Office clients. The entity model simplifies application maintenance since most changes made on external data sources only require updates in the entity (that is, the ECT) rather than a change to every application or site that uses it.

BCS Delivered in Multiple Products

The components of BCS are distributed among multiple products (for an illustration, see "BCS Component Distribution"), as follows:

SharePoint Foundation 2010 includes the core run-time service for BCS plus components that allow external data to be accessed and displayed in SharePoint lists.

SharePoint Server 2010 includes components for searching external sources and mapping credentials to external systems, reusable Web Parts to provide read-only access, and extensions to allow users to access data within Office 2010 applications.

Office 2010 includes components that allow access through SharePoint Server 2010 to external data sources, with user interfaces that vary for each application.

BCS will be used as the data access technology for the next version of Duet, a joint Microsoft-SAP product that enables employees to view and enter SAP data in Outlook and other Office applications, rather than an SAP client. For example, workers might enter time worked or vacation requests in the Outlook calendar. The upcoming version, Duet Enterprise for Microsoft SharePoint and SAP, is expected in the second half of 2010. Also, the search functionality in SharePoint and Search Server uses BCS connectors to crawl and index data in external sources.

Development Tools Span Skill Levels

BCS capabilities can be implemented using SharePoint Designer or VS 2010. SharePoint Designer lets nondevelopers design and customize SharePoint sites through a forms-based user interface without writing code. Simple ECTs can be implemented and deployed with SharePoint Designer without requiring a skilled developer. For example, an end user in a sales department who has basic connection information and adequate permissions on products and sales tables in a SQL Server database could use SharePoint Designer to create an ECT that exposes that data in SharePoint lists.

VS 2010 offers more advanced tools than SharePoint Designer to implement solutions that use BCS. A new graphical designer for ECTs and full support of the VS debugger is included. Complex ECTs that require aggregating data from multiple sources, transforming data, using business logic procedures, or custom security protocols must be developed with the VS 2010 tools. For example, an ECT that cross-references sales data from a SQL Server database with customer data from a CRM solution would be created by a developer using VS 2010. ECTs originally created using SharePoint Designer can be opened in VS 2010 as a starting point to create a more complex ECT.

ECTs can be deployed to the BCS metabase on a SharePoint server from within either VS 2010 or SharePoint Designer. VS 2010 allows an ECT to be packaged into a deployment file that can be delivered to a SharePoint server and deployed with minimal manual steps.

Client Choices Offer Flexibility

External data served by BCS can be accessed by end users through SharePoint's Web-based client or Office 2010 applications, including Word, Outlook, Excel, Access, InfoPath, and SharePoint Workspace. Organizations with varied data solutions often have small groups of users who have the skills and knowledge to retrieve data from specific server sources. Consequently, these users get inundated with requests from other users to access and update such data. SharePoint and BCS may improve this situation by granting simple data access to a broader set of users within common client applications such as Word and Outlook. For example, client contacts stored on a third-party CRM database could be exposed in an Outlook contacts view using Office 2010, SharePoint 2010, and BCS.

Using an Office 2010 application as a BCS client requires SharePoint Server 2010 (SharePoint Foundation 2010 cannot be used by itself to support BCS in Office 2010), and each accessing user will need a SharePoint Server 2010 Enterprise Client Access License (CAL). User interface elements that access BCS from within Office applications, such as forms and buttons, are delivered to the client as Visual Studio for Office add-ins using the ClickOnce packaging and deployment technology.

An offline cache allows Office 2010 users to download data from an external source while online and work with that data while disconnected from the network. When the client reconnects, changes made to the data can be synchronized with the servers. For example, in an Outlook-based field service client, a user could enter information about a service call in Outlook while offline and update the field service application when back online.

Developers Still Needed

Although the out-of-the-box capabilities of BCS provide much more infrastructure and higher-level implementation tools than were provided with the SharePoint 2007 BDC, many organizations will find that BCS still requires custom code written by skilled developers.

Back-end data sources often contain sensitive, business-critical data that demands carefully controlled access, especially when changes are made. BCS provides the hooks to incorporate data filters and business logic that can provide end users with appropriate access that protects back-end data, but custom code will be required in most cases.

SharePoint Tools in Visual Studio 2010

New tools for SharePoint application development are built in to Visual Studio (VS) 2010. The tools and templates could make SharePoint application development simpler and more familiar to .NET developers, which will help partners and corporations customize the product.

Building on Windows SharePoint Services Extensions

The VS 2010 features for SharePoint build on the free Visual Studio extensions for Windows SharePoint Services, also known as VSeWSS.

SharePoint applications are based on ASP.NET 2.0 and the Windows Workflow Foundation, so VS 2008 can be used for SharePoint development even without the extensions for WSS. However, the extensions enhance VS 2008 with development support for SharePoint site functionality, such as specialized Web Parts, lists, content types, and event receivers (custom code that is run when events occur on a site's server). The extensions for WSS also include a tool that displays the contents of SharePoint Solution package files used to deploy SharePoint sites, and a stand-alone tool that generates a VS solution file from an existing SharePoint site. Version 1.3 of the extensions, which was released in a Community Technology Preview version in Mar. 2009 but has not been offered as a final release, is the last upgrade, as SharePoint integration is now built into VS 2010.

VS 2010 includes all capabilities of the VS extensions for WSS, and more.

New Server Exploration, Visual Designers, Deployment Tools

VS 2010 extends the Server Explorer window to connect to SharePoint 2010 and 2007 servers and display SharePoint application components, such as content types, lists, surveys, and workflows. Component detail is displayed in the VS Properties window.

A visual Web Part designer makes it easier to create and modify Web Parts within VS 2010. A wizard interface simplifies event receiver creation, and a visual designer is provided for designing forms used to request workflow information.

VS 2010 allows a SharePoint solution package file to be imported into a VS project without using the separate stand-alone tool required by the extensions for WSS. An improved design tool eases manipulation of deployment packages, and the client and server portion of an application can be included in the same package. Deployment from VS to a SharePoint server is completely automated so that a single keystroke builds, packages, and deploys the application to the server in one step.

All project templates available with the VS extensions for WSS are included in VS 2010 and a Business Data Connectivity Model template is added for development with SharePoint Server 2010. Business Connectivity Services is a component of SharePoint Server 2010 that acts as a bridge between the Web portal and back-end databases and business applications, such as customer relationship management or enterprise resource planning systems.

Development Teams May Benefit

SharePoint application development teams that use VS with Team Foundation Server (TFS) may realize further benefits. The SharePoint integration in VS 2010 opens SharePoint applications to features such as source control, bug tracking, and build automation when TFS is used.

SharePoint Server 2010 is the fourth major release of the product since it was derived from Site Server in 2001. So far, an inconsistent, loosely coupled development experience has hindered SharePoint in the development community. The out-of-the-box SharePoint integration in VS 2010 may improve SharePoint's reputation in this regard by simplifying SharePoint application development.

Other New SharePoint 2010 Development Capabilities

SharePoint Foundation 2010 provides several new and enhanced development capabilities in addition to Business Connectivity Services and the Visual Studio tools, including the following:

Sandboxed applications. SharePoint Foundation 2010 allows custom applications to run in a sandboxed mode that limits use of resources (such as processor time) and restricts access to data (such as the file system), reducing the damage a runaway application might cause. A sandboxed application that surpasses its resource quota is automatically shut down. This will reduce the risks associated with SharePoint customization and will particularly benefit SharePoint hosting companies (including Microsoft itself) that want to allow customization by their tenants.

LINQ to SharePoint. LINQ is a set of APIs and programming language features that enables data access queries to be written in .NET programming languages such as C# rather than as text data that is not parsed by the code editor. SharePoint Foundation 2010 delivers LINQ to SharePoint, which provides IntelliSense command completion and strong typing of identifiers within SharePoint list queries to help catch coding errors while writing and compiling code rather than at run time.

Client object model. SharePoint Foundation 2010 provides new client APIs that allow access to SharePoint data and services from remote computers running .NET-based, Silverlight-based, and JavaScript-based applications. These APIs will particularly benefit organizations that use custom client-side code to read or write SharePoint objects without running specialized code on the server.

Developer dashboard. Developers can view statistics showing the resources used by a SharePoint server when it generates a Web page, which can help identify bottlenecks in custom code and locate inefficient code that could overburden a server.

Workflow improvements. The workflow capabilities of SharePoint Foundation have received several improvements. New activities, such as activities that ease date and string manipulations, are provided. A new type of workflow called a pluggable workflow service can await a response from an external service (from an outside vendor, for example) before proceeding. Workflows can now generate events, enabling external code to monitor a workflow and take action in response; among other things, this enables developers to write code that extends the logic of built-in SharePoint workflows. Workflows can now be associated with a SharePoint site and started by a site user with the appropriate permissions rather than being triggered in response to events on a specific SharePoint list item, as was required in the previous SharePoint release. Declarative workflows can now be reused, allowing a workflow to be created once and used for multiple applications.

Sync Framework integration. SharePoint Foundation 2010 supports synchronizing lists and document libraries between a SharePoint server and a client application using the services and architecture of the Microsoft Sync Framework, a platform for synchronizing data between independent databases, file systems, and other data sources. This Sync Framework support could allow other applications that work with the Framework to synchronize data with SharePoint without additional custom code.

InfoPath 2010 and SharePoint Server 2010 also deliver improvements for designing SharePoint forms.

Impact and Roadmap

Overall, SharePoint 2010 is a major release that delivers many improvements. Even organizations that don't plan an immediate upgrade might need to evaluate SharePoint 2010 in connection with other products released around the same time. A new SharePoint version is unlikely to appear until 2013.

SharePoint 2010 Affects Other Decisions

As described above, SharePoint 2010 integrates with many other products, including Office 2007 and 2010, SQL Server 2008 R2, and Project Server 2010. Organizations can evaluate the latest versions of each of these other products in isolation, as each delivers benefits over its predecessors. However, many of the most interesting new capabilities require SharePoint 2010. (See the illustration "2010 Product Relationships".) Consequently, the decision to adopt one product in the 2010 generation can commit an organization to adopt several others.

For example, the Office Web Apps and coauthoring features for group work on Office documents require SharePoint 2010. Similarly, the PowerPivot feature requires Excel 2010, SharePoint Server 2010, and SQL Server 2008 R2. SharePoint 2010 in turn requires 64-bit editions of SQL Server (2005 SP3 or later) and Windows Server (2008 or later). SharePoint 2010 sites work with Internet Explorer (IE) 7 (not IE 6), Firefox 3, and Safari 3 and later versions of these three browsers, although IE and Firefox on Windows will have some features not available on other OSs.

Organizations can start by evaluating each product's new version in isolation, taking into account which version they have already deployed and what they have already licensed through programs such as Software Assurance. However, organizations that do want to evaluate the products together might want to start with SharePoint 2010: it is required for new Office 2010 collaboration features and the Office Web Apps, as well as the new Visio Services feature for distributing diagram-based reports to browser users. SharePoint Server 2010 also helps centrally distribute and manage spreadsheets created with PowerPivot, although it is not required. Finally, SharePoint Server 2010 is a prerequisite for Project Server 2010. Therefore, organizations evaluating the 2010 releases might want to decide on their SharePoint 2010 plans first, and then make plans for other releases with SharePoint in mind.

Beyond SharePoint 2010

Microsoft has not announced software releases beyond SharePoint 2010, but it is possible to speculate based on past release patterns and current features of the product. (See the illustration "SharePoint Roadmap Overview".)

First, the company plans to update its SharePoint Online Dedicated and Standard hosted services to use SharePoint 2010 in the second half of 2010. Both SharePoint 2010 Online services will probably introduce capabilities, such as records management, that were previously only available with the on-premises software. The updated Dedicated service might also introduce some capabilities not available in the Standard service, such as FAST Search for SharePoint. The Dedicated service could also offer SharePoint 2010 earlier than the Standard service does, as has happened with Exchange Server 2010.

The next full SharePoint software version (labeled SharePoint 2013 here) will probably appear in late 2013 at the earliest. The company has not announced any priorities, but possible ones might include the following:

Document sharing and collaboration. SharePoint will probably offer an out-of-the-box option for storing documents and other large content in file shares and other remote storage devices, building on the new SharePoint 2010 remote storage APIs. This and other changes will also probably enable Microsoft to raise SharePoint's maximum file size, currently 2GB. Office Web Apps will also probably add more capabilities from their corresponding products in the Office suite, such as revision tracking in Word and chart creation in Excel.

Search. All of the SharePoint and Search Server products will probably get a new search platform based on the just-released FAST Search Server for SharePoint, with improvements to simplify setup, customization, and management.

Content management. SharePoint's features for managing content retention and destruction will probably be extended to more SharePoint data (such as user status updates, shared bookmarks, and other social data) and get improvements for managing content from other sources (such as file shares and Exchange mailboxes).

Business intelligence. A future SharePoint version might consolidate some overlapping BI technologies; for example, it might deliver a single charting component in place of the three separate ones used by PerformancePoint Services, Excel Services, and the Chart Web Part. More radically, the company might enhance the key performance indicator, scorecard, and dashboard capabilities of Excel and Excel Services to take over for PerformancePoint Services. Separately, the PowerPivot analytics plug-in for SharePoint might get improvements to handle larger spreadsheets (it is currently limited to 2GB) and to help upsize solutions from PowerPivot to SQL Server Analysis Services.

Scalability, availability, and management. SharePoint will probably get continued improvements for large-scale hosting, such as more complete multitenant management. Replication and other improvements are likely to help geographically distributed organizations manage multiple SharePoint farms. The product's Central Administration console will probably gain the ability to record PowerShell scripts to capture administrative procedures, something already supported by other Microsoft server products.

Development. The Business Connectivity Services application integration technology will probably get continued improvements, driven in part by Microsoft's own use of the technology for products such as Dynamics CRM and the Duet Enterprise connectivity product for SAP applications.

Appendix: SharePoint Server 2010 Packaging, Pricing, and Licensing

Packaging, licensing, and pricing changes introduced with SharePoint Server 2010 in May 2010 include moving developer-related components from the commercial product into the free SharePoint Foundation offering, a 10% price increase for certain server and Client Access Licenses (CALs), a new server license covering use of higher-end search technology than is included in the base product, and an additional server license for Internet and extranet sites. While the changes will probably increase Microsoft's revenue over comparable license sales for SharePoint 2007, customers will also get more value.

Retains Similar Packaging and Licensing Model

The SharePoint 2010 product line offers four tiers of functionality, up from three in SharePoint 2007, with each tier being a superset of the one below. SharePoint is licensed using a server-CAL model, under which an organization purchases a server license for the right to run software on a server, and CALs for the right for users or devices to access the servers.

The four main tiers are as follows:

SharePoint Foundation. SharePoint Foundation 2010 is a free download that can be deployed by customers who have already purchased Windows Server 2008/2008 R2 server licenses and Windows Server 2008 CALs. Formerly called Windows SharePoint Services, it includes tools and services for creating and managing team collaboration sites and provides basic services on which SharePoint Server 2010 relies. (See the sidebar "SharePoint Foundation 2010".)

SharePoint Server with Standard CAL. The next tier of functionality is licensed by purchasing SharePoint Server 2010 server licenses along with a Standard CAL for each client. This level licenses SharePoint Server's core team collaboration, corporate portal hosting, enterprise search, Web content management, and document management features.

SharePoint Server with Standard and Enterprise CALs. The addition of the Enterprise CAL to SharePoint Server and the Standard CAL provides the right to use all SharePoint Server 2010 BI tools, as well as components that allow nonprogrammers to create certain types of business collaboration solutions without having to write custom code.

SharePoint Server with Standard CAL, Enterprise CAL, and FAST Search Server for SharePoint. Like its predecessors, SharePoint Server 2010 includes indexing and query capabilities for searching Intranet content. (Customers who want only these search capabilities can license a special-purpose edition of SharePoint called Search Server 2010; see the sidebar "Search Server 2010".) However, customers can get better search tuning, greater scale, and other improvements with FAST

Search Server for SharePoint (FAST for SharePoint). FAST for SharePoint works only with SharePoint Server 2010 and requires SharePoint Server Standard and Enterprise CALs. It also requires its own server licenses—which are four times as expensive as SharePoint Server 2010 server licenses—for the subset of servers that implement the search farm.

For scenarios involving Internet or extranet sites accessed by nonemployees, Microsoft continues to offer an alternative to the server-CAL model described above; customers can choose whichever model is less expensive or more practical to manage. Under the alternative model, SharePoint Server 2010 for Internet Sites Standard licenses an instance of SharePoint that can be accessed by an unlimited number of nonemployee clients who utilize SharePoint Standard CAL-level features, while SharePoint Server 2010 for Internet Sites Enterprise licenses nonemployees for the full SharePoint feature set as well as use of FAST for SharePoint. The SharePoint Server 2010 for Internet Sites Enterprise license is similar to a predecessor, the SharePoint Server 2007 for Internet Sites license, while the Standard license is new.

(For a brief description and price for the dozen licenses associated with running SharePoint Server 2010 on-premises, see the chart "Pricing for SharePoint-Related Licenses". All prices quoted in this article represent the highest price a U.S. business customer would pay through the Open No Level, formerly called Open Business, volume purchasing program.)

Server-Side Licensing

Besides requiring licensing for SharePoint Server, all servers in an organization's internal SharePoint deployment must be licensed for Windows Server, and one or more servers will need to be licensed for SQL Server and possibly others for FAST for SharePoint.

SharePoint Server

The full SharePoint Server 2010 product comes in only one edition for internal use. This is unlike other Microsoft servers licensed with the server-CAL model—Windows Server, Exchange, SQL Server, and Communications Server offer at least two editions, Standard and Enterprise, which generally differ with respect to IT-focused capabilities such as scalability and fault tolerance (rather than end-user features).

The software associated with the SharePoint Server 2010 server license (US$4,926) contains the code to implement both the Standard CAL-level features as well as the Enterprise CAL-level features. During product configuration, the administrator must explicitly choose to enable Enterprise CAL-level capabilities. This is done for license compliance purposes and is mainly an attempt to prevent organizations from inadvertently obliging themselves to purchase Enterprise CALs.

Note that some Microsoft documents use the terms "SharePoint Server 2010 Standard Edition" and "SharePoint Server 2010 Enterprise Edition," but this somewhat misleading terminology does not refer to two separate server license types. Rather, "SharePoint Server 2010 Standard Edition" refers to a server licensed for SharePoint Server 2010 that has the Standard CAL

features enabled, but not the Enterprise CAL features. And "SharePoint Server 2010 Enterprise Edition" refers to a server licensed for SharePoint Server 2010 that has all (both the Standard CAL and Enterprise CAL) features enabled.

SharePoint Server contains several discrete components, but a SharePoint Server 2010 license is required no matter how much of the software is installed and running. (The exception is Search Server, which is a special-purpose subset of SharePoint Server and is licensed separately.) Small and midsize organizations not requiring fault tolerance could install all the SharePoint components on a single server, but larger organizations generally use several servers, often dedicating each server to a particular set of SharePoint tasks, for reasons of performance, high availability, administrative flexibility, and security. In both scenarios, each server running any SharePoint Server 2010 component or set of components requires a license. (For an overview of rules governing SharePoint Server 2010 server license reassignment and number of running instances per license, see the sidebar "SharePoint Server Virtualization Rules".)

FAST for SharePoint

While SharePoint Server 2010 includes crawling, indexing, and query capabilities for searching Intranet content, Microsoft also provides customers with a higher-end option called FAST for SharePoint. (See the chapter "Enterprise Search in SharePoint 2010".)

Customers deploying FAST for SharePoint would almost certainly have a multiserver SharePoint implementation with one or more servers devoted to it.

Each server that runs FAST for SharePoint components must be licensed with a FAST for SharePoint server license (US$21,975). If the server is dedicated to FAST for SharePoint—which is the most likely configuration for production scenarios—a SharePoint Server 2010 server license is not required for that server. However, FAST for SharePoint requires at least one instance of SharePoint Server 2010 to be deployed in the same farm, and those servers must be licensed appropriately.

SQL Server

SQL Server is required to store SharePoint configuration parameters, security credentials, content, usage logs, web analytics reports, state information, and other items. Customers have both free and fee-based options.

SQL Server 2008 (or 2008 R2) Express edition. Called Microsoft SQL Desktop Engine (MSDE) in earlier releases, the free Express edition can be used as a database back-end but is practical only for small-scale SharePoint sites. Express can exploit only a single processor and no more than 1GB of RAM, and limits database size to either 4GB or 10GB (depending on version). This limits the number of documents that can be stored and indexed and the amount of health monitoring and usage data that can be maintained.

SQL Server 2005, 2008, or 2008 R2 Standard or Enterprise editions, or 2008 R2 Datacenter. Fee-based editions of SQL Server are able to exploit more memory, processors, and much larger databases. While SQL Server Standard should be sufficient in most cases, a few SharePoint features—such as crawling (for search indexing) and Web analytics reporting—can benefit from data compression, table partitioning, and other features present in Enterprise and Datacenter, but not Standard. (Table partitioning improves query performance on large database tables that are split across multiple disk systems or database servers.) SQL Server Workgroup and Web editions are not supported for use with SharePoint Server 2010.

The Standard and Enterprise editions can be licensed under a server-CAL or per-processor model, while Datacenter is licensed exclusively per-processor.

Windows Server

Every server that runs SharePoint Server or FAST for SharePoint components must be licensed for Windows Server 2008 or 2008 R2 Standard, Enterprise, or Datacenter edition. Editions of SQL Server 2005, 2008, and 2008 R2 can run on versions of Windows Server going back to Windows Server 2003.

The biggest factor determining which Windows Server edition to use is the OS's memory support. Windows Server Standard supports up to 32GB, which should be sufficient for most SharePoint-related work roles, and Enterprise up to 2TB.

Client-Side Licensing

Employee access to SharePoint servers is licensed through the purchase of SharePoint Server CALs. (For information about the one exception, see the sidebar "Project Server 2010 and SharePoint".)

Like CALs for most other Microsoft server products, SharePoint CALs must be assigned to each user or each accessing device, usually a PC. Each SharePoint client also needs a Windows Server 2008 (or later) CAL, as well as a SQL Server 2005 (or later) CAL if SQL Server is licensed in server-CAL mode.

As with SharePoint Server 2007, accessing the full SharePoint Server 2010 feature set requires customers to purchase two different CALs, a SharePoint Server 2010 Standard CAL (SCAL) and an Enterprise CAL (ECAL). The price for the SharePoint 2010 SCAL remains at US$95, and the ECAL is US$83, up 10% from SharePoint Server 2007.

The SCAL, always required, licenses client access to SharePoint's team collaboration, corporate portal hosting, enterprise search, Web content management, and document management features. An ECAL adds rights to use all of SharePoint Server 2010's BI tools, useful for summarizing data collected in a company's day-to-day operations to help workers make business decisions. It also licenses use of two sets of out-of-the-box SharePoint components and tools—Access Services and InfoPath Forms Services—that make it possible for nonprogrammers to create business collaboration solutions without having to write custom code. Approximately half the capabilities licensed by the ECAL are new to SharePoint Server 2010.

An ECAL is required to use any of the features described in the following sections.

Excel Services

Excel Services makes it possible for browser users to view an Excel workbook stored on SharePoint Server and temporarily change input values in the workbook to perform "what if" analysis. In SharePoint 2010, Excel Services can also use PowerPivot for SharePoint to process large multidimensional databases (sometimes called cubes) embedded in Excel workbooks. (See the chapter "Improved Dashboards, Reports from SharePoint".) Each client accessing workbooks published through Excel Services in SharePoint Server 2010 requires a SharePoint 2010 ECAL. Publishing a workbook to Excel Services requires Excel 2007 or 2010. Any server running the PowerPivot for SharePoint components must be licensed for SQL Server 2008 R2 Enterprise or Datacenter—key components of PowerPivot for SharePoint are delivered in SQL Server 2008 R2 Enterprise (and Datacenter) media and installed from those media on SharePoint servers.

Excel Services shares some of the same technology used in the Excel Web App add-on, which enables SharePoint 2010 users to create, edit, and view Excel workbooks using only a browser. (See the chapter "Document Sharing with Office 2010 and SharePoint 2010".) When run on an on-premises SharePoint server, Excel Web App, like the other Office Web Apps, requires every user to be licensed for Office 2010 Standard or Professional Plus, but it does not require the SharePoint ECAL (and can run on the free SharePoint Foundation). Excel Services, in contrast, does not require an Office license except for the users who publish workbooks, but it does require a SharePoint ECAL for all users.

Visio Services

New to SharePoint Server 2010, Visio Services makes it possible for browser users to view Visio diagrams stored on and rendered by SharePoint Server 2010 without having Visio installed on the client. (See the chapter "Improved Dashboards, Reports from SharePoint".) Each client accessing diagrams published through Visio Services requires a SharePoint 2010 ECAL. Publishing diagrams via Visio Services requires that a user create diagrams in the new Visio Graphics Service (VDW) format—a capability exclusive to Visio 2010 Professional (US$459) and the new Visio 2010 Premium (US$810). No edition of Visio is included in any Office 2010 suite.

Access Services

New to SharePoint Server 2010, Access Services gives SharePoint Server 2010 the ability to host applications created using Access 2010 and make them accessible to users from a browser, greatly simplifying both the initial rollout of a database application as well as the deployment of subsequent modifications. (See the chapter "Improved Dashboards, Reports from SharePoint".)

Each client making use of Access Services requires a SharePoint 2010 ECAL. All design work occurs in Access 2010—anyone wanting to publish an application to Access Services (or modify an existing application published to Access Services) requires Access 2010, which is available separately or as part of Office 2010 Professional Plus, but not Office 2010 Standard.

PerformancePoint Services

SharePoint Server 2010's PerformancePoint Services helps organizations track business performance and health and flag potential problems. (See the chapter "Improved Dashboards, Reports from SharePoint".) Each client accessing pages published through PerformancePoint Services requires a SharePoint 2010 ECAL.

InfoPath Forms Services

InfoPath Forms Services makes it possible for browser users who do not have the InfoPath forms-building client to fill out forms that were created with that client.

InfoPath Forms Services was first introduced in SharePoint Server 2007 and required a SharePoint 2007 ECAL. (It was also available in a separate Forms Server product that was discontinued in 2009.) As in earlier SharePoint versions, each client making use of browser-based forms presented by InfoPath Forms Services in SharePoint 2010 requires a SharePoint 2010 ECAL. All forms design work occurs in InfoPath 2010. InfoPath 2010 is available separately, or in Office 2010 Professional Plus, but not Office 2010 Standard.

FAST for SharePoint

Each client submitting queries to the FAST for SharePoint engine (described above) requires a SharePoint 2010 ECAL, and the servers hosting FAST Search components also require server licenses as described above.

Business Connectivity Services

First introduced in SharePoint Server 2007 as the Business Data Catalog (BDC) and renamed and improved in SharePoint Server 2010, Business Connectivity Services (BCS) enables users to view, update, and search selected data in sources external to SharePoint, such as CRM and enterprise resource planning systems. (See the chapter "SharePoint 2010 Improves Tools, APIs".)

The BDC technology stack shipped with SharePoint Server 2007 and any use of the BDC triggered the need for the SharePoint 2007 ECAL. However, with the latest release, SharePoint Foundation 2010 includes the core run-time service for BCS as well as components that allow external data to be accessed and displayed in SharePoint lists.

SharePoint Server 2010 adds the components necessary for the following features:

- Indexing external sources so that they can be searched through SharePoint's enterprise search capabilities or FAST Search (this feature is also provided with Search Server, which requires no CALs)
- Mapping of account credentials for connecting to external systems (i.e., a SharePoint user account can be mapped to a CRM application account so users get appropriate access through SharePoint)
- Displaying information extracted from external systems within a browser (i.e., supporting Web Parts)
- Accessing external data from within Office 2010 applications via SharePoint Server.

Each client that makes use of any of the BCS-related components shipped with SharePoint Server 2010 requires an ECAL. This includes any client submitting search queries to a SharePoint system that indexes external content via the BCS, and any client viewing pages consisting of Web Parts that display data extracted via the BCS.

Chart Web Part

The SharePoint 2010 Chart Web Part provides a relatively simple way to incorporate charts into SharePoint pages. (See the chapter "Improved Dashboards, Reports from SharePoint".)

Any client building or viewing pages containing a Chart Web Part requires an SharePoint 2010 ECAL. (Note that Microsoft's quarterly Product Use Rights document uses the more generic term "Advanced Charting" rather than "Chart Web Part" in its list of features requiring an ECAL.)

Licensing for Internet Use

Customers using SharePoint Server to construct Internet or extranet sites accessed by nonemployees—such as customers or business partners—get the option to choose between two licensing approaches. While customers can use the server-CAL approach, it may not be practical for the organization to identify the set of users or devices that would require CALs, or the set of possible users could be so large as to make CAL purchases prohibitively expensive. To address such situations, Microsoft offers an alternative specifically for licensing nonemployee users.

While most other Microsoft server products licensed under the server-CAL model provide an alternative to CALs for licensing nonemployee users—usually in the form of an External Connector license—SharePoint's alternative model is unique. Rather than require customers to buy two different server licenses for each server—a license for the right to run the server software and a separate External Connector license to allow an unlimited number of clients to access the software running on the server—SharePoint bundles both these rights into a single license.

A SharePoint Server for Internet Sites server license includes both the right to run the SharePoint software on a server and the right for an unlimited number of nonemployee clients (on-site contractors are considered employees) to access the server. The license also covers server access by employees (i.e., employees wouldn't need CALs), but only if either of the following conditions are true:

- The SharePoint server makes exactly the same content available to employees and nonemployee users alike (in other words, if employees can access content, nonemployees must be able to access it as well)
- The employee is accessing the site for the sole purpose of authoring content for external users.

In multiserver deployments licensed using SharePoint Server for Internet Sites server licenses, all SharePoint servers that directly or indirectly provide content to nonemployee users require a SharePoint Server for Internet Sites license. As one might expect, this includes SharePoint front-end servers and application servers. However, it also includes indexing servers and even servers used for content staging. Servers in a development or test environment—for example, used to customize the look of the Web site and test it before changes go into production—are an exception, as they are typically covered by Microsoft Developer Network (MSDN) licensing.

Types of Licenses for Internet Sites

There are two types of SharePoint Server for Internet Sites licenses.

SharePoint Server 2010 for Internet Sites, Standard (US$11,793) provides the right to run a single instance of SharePoint Server 2010 on a server with only Standard CAL-level capabilities enabled. The software may be used to create and host a single site at a time resolved from a single domain name (e.g., directionsonmicrosoft.com) and its subdomains (e.g. licensing.directionsonmicrosoft.com). Microsoft defines a site as a collection of interconnected Web pages and related Web content that share a common home page. This license is new for SharePoint Server 2010.

SharePoint Server 2010 for Internet Sites, Enterprise (US$41,392) provides the right to run (on the licensed server) either SharePoint Server 2010 with both SCAL-level and ECAL-level capabilities enabled, or FAST for SharePoint. The terms of the license allow a single instance of SharePoint Server 2010 or FAST for SharePoint, but not both, meaning an external-facing SharePoint site that uses FAST for SharePoint will require at least two licenses for SharePoint Server 2010 for Internet Sites, Enterprise. Note that all servers comprising a site that exposes ECAL-level features to external users require a SharePoint Server 2010 for Internet Sites, Enterprise license—even servers that do not host SharePoint components that directly support ECAL-level capabilities. Customers with an existing SharePoint Server 2007 for Internet Sites license covered under Software Assurance (SA) as of Apr. 30, 2010, receive upgrade rights to SharePoint Server 2010 Internet Sites, Enterprise.

The software associated with both the Standard and Enterprise edition licenses are installed from the same media as SharePoint Server 2010. The installer can differentiate between a license key for Standard and Enterprise and activate the appropriate set of features. Neither license permits the hosting of Office Web Apps.

In most cases, besides licensing each server with an edition of SharePoint Server 2010 for Internet Sites, customers also have to license each server with a Windows Server 2008 External Connector (US$2,020), which provides the right for an unlimited number of nonemployee clients to access all instances of Windows Server running on a physical server. (The most common exception to this rule: if users are not authenticated or otherwise individually identified in any way by the SharePoint site, no

Windows Server External Connector is required.) Furthermore, one or more SQL Server per-processor licenses are required for the SharePoint site's back-end database.

Sites Used by Both External and Internal Users

Customers who want to use the same SharePoint Server or group of servers to support both employee-specific content as well as externally accessible content have to double-license each server instance, applying both a SharePoint Server 2010 server license and a SharePoint Server 2010 for Internet Sites (Standard or Enterprise) license to each running instance of the server software. As with internal-only sites, employee access requires the purchase of SharePoint Server CALs.

Similarly, customers wanting to serve both internal and external users with a single FAST for SharePoint deployment may acquire licenses for both FAST for SharePoint and SharePoint Server 2010 for Internet Sites, Enterprise and assign those licenses to the same instance running on a server.

License Compliance Tool on the Way?

At the Oct. 2009 SharePoint conference, Microsoft indicated the company was working on a license usage report tool to help customers measure the number of SharePoint 2010 SCALs and ECALs that would be required for license compliance given the way SharePoint was being used within their enterprise. Among other things, the tool was advertised as tracking the number of unique users using SCAL features only, using ECAL features, and using Office Web Apps. However, the tool did not ship with SharePoint Server 2010 in May, and presumably this tool will be delivered later this year as a download.

Resources

General

A SharePoint 2010 planning and architecture guide is at technet.microsoft.com/library/cc261834(office.14).aspx.

For background on SharePoint and past SharePoint versions, see the Apr. 2007 *Research Report,* "SharePoint Platform Matures, Expands Role," and the Sept. 2003 *Research Report,* "Collaboration and Portal Strategy Built on SharePoint."

Office Document Sharing

For an evaluation guide to Office 2010 as a standalone suite, see the May 2010 *Research Report,* "Office 2010 on the Desktop."

Product guides for the Office Web Apps and Office 2010 client applications are available at www.microsoft.com/downloads/details.aspx?FamilyID=e690baf0-9b9a-4c47-88da-3a84f3e9b247.

Deployment information for Office Web Apps is available at technet.microsoft.com/en-us/library/ee695756(office.14).aspx.

A white paper with deployment information for PowerPoint's Broadcast Slide Show feature can be downloaded at www.microsoft.com/downloads/details.aspx?FamilyID=3a8ff829-d0e6-4b35-b637-16780d238fec.

The File Synchronization Service via SOAP over HTTP (MS-FSSHTTP), a new Microsoft protocol used to exchange data in coauthoring scenarios, is documented at msdn.microsoft.com/en-us/library/dd943623.aspx.

Web Collaboration

SharePoint Server 2007 Web collaboration features and their use are outlined at sharepoint.microsoft.com/social.

SharePoint Server 2010 privacy controls for users and administrators are summarized at technet.microsoft.com/library/ee620541.aspx.

The Outlook Social Connector is described in more detail at blogs.msdn.com/outlook/archive/2009/11/18/announcing-the-outlook-social-connector.aspx.

Enterprise Search

An overview of enterprise search improvements in SharePoint Server 2010 is available at technet.microsoft.com/library/ee667266%28office.14%29.aspx.

An outline of the differences between the four SharePoint enterprise search products discussed in this article is available at www.microsoft.com/downloads/details.aspx?familyid=D7C0091E-5766-496D-A5FE-94BEA52C4B15.

A 62-page evaluation guide for FAST Search Server 2010 for SharePoint is available at www.microsoft.com/downloads/details.aspx?FamilyID=c422d3c7-1443-41e4-b0fe-fc402ee4d8c1.

TechNet documentation for FAST Search Server 2010 for SharePoint begins at technet.microsoft.com/library/ee781286%28office.14%29.aspx.

TechNet documentation for Search Server Express begins at technet.microsoft.com/library/dd183108%28office.14%29.aspx.

A diagram explaining planning considerations for SharePoint 2010 search architecture is available at www.microsoft.com/downloads/details.aspx?familyid=5655EACA-22DF-4089-BCD3-38A1F5318140.

A diagram with examples of various SharePoint 2010 search architectures and guidance for tuning aspects of enterprise search performance are available at www.microsoft.com/downloads/details.aspx?familyid=22FFC029-2C08-457D-8311-CA457C6D160E.

A diagram of some farm-level enterprise search architectures is at www.microsoft.com/downloads/details.aspx?familyid=5A3CA177-FB9A-4901-9797-0C384277DB7C.

Hardware requirements for SharePoint 2010, including a link to download the Speech Platform component required for phonetic name-matching in people search, are at technet.microsoft.com/library/cc262485%28office.14%29.aspx.

Content Management

Retention policy enforcement in Exchange Server 2010 is explained in the Jan. 2010 *Research Report,* "Evaluating Exchange Server 2010."

Windows Server's File Classification Infrastructure is discussed in more detail in the Dec. 2009 *Research Report,* "Evaluating Windows Server 2008 R2."

EMC content management products and SharePoint integration are explained in more detail at www.emc.com/campaign/global/captiva/index.htm#/page/get-paper-into-sharepoint/ (Captiva), www.emc.com/products/detail/software/my-documentum-microsoft-sharepoint.htm (My Documentum), and www.emc.com/products/detail/software/sourceone-microsoft-sharepoint.htm (SourceOne).

Word Automation Services is summarized at blogs.msdn.com/microsoft_office_word/archive/2009/10/26/introducing-word-automation-services.aspx and msdn.microsoft.com/library/ee558830(office.14).aspx.

Business Intelligence

PowerPivot is outlined in "PowerPivot Strengthens Analytics in Excel, SharePoint 2010" on page 21 of the May 2010 *Update*.

Visio Services is described in more detail in "Visio 2010 Gains Ribbon, SharePoint Service" on page 10 of the May 2010 *Update*.

A Microsoft SharePoint BI site is technet.microsoft.com/sharepoint/ee692578.aspx.

A SharePoint 2010 BI planning guide appears at technet.microsoft.com/library/ee683867.aspx.

Basic criteria for choosing among SharePoint BI technologies appear on a Microsoft poster at go.microsoft.com/fwlink/?LinkId=167170.

The end for PerformancePoint Server 2007 is explained in more detail in "PerformancePoint Serviced" on page 12 of the Dec. 2009 *Update*.

SQL Server 2008 improvements are outlined in the Dec. 2008 *Research Report,* "SQL Server 2008 Enhances Management, Business Intelligence."

Reporting Services backward compatibility in SQL Server 2008 and 2008 R2 is discussed in more detail at msdn.microsoft.com/library/ms143251.aspx.

For an overview of Microsoft's BI strategy as of 2007, see the May 2007 *Research Report,* "Microsoft's Business Intelligence Strategy."

Scalability, Availability, Management

Administrative planning and procedures documents are available for SharePoint Foundation 2010 at technet.microsoft.com/library/cc288070.aspx and for SharePoint Server 2010 at technet.microsoft.com/library/cc303422.aspx.

A list of major SharePoint partners appears at sharepoint.microsoft.com/businessproductivity/resources/pages/partners.aspx.

EMC Documentum Repository Services for SharePoint is described in more detail at www.emc.com/products/detail/software/documentum-repository-services-sharepoint.htm.

Windows SharePoint Services 3.0 SP2 and SharePoint Server 2007 SP2 information and downloads are available via support.microsoft.com/kb/968170.

Development

Microsoft's SharePoint Developer Center, which links to developer documentation and videos, is at msdn.microsoft.com/sharepoint/default.aspx.

Guidance for migrating SharePoint 2007 Business Data Catalog solutions to the SharePoint 2010 Business Connectivity Services is at technet.microsoft.com/library/ff607947.aspx.

More information about Duet is at www.duet.com.

New capabilities of InfoPath 2010 are summarized at blogs.msdn.com/b/infopath/archive/2009/07/15/what-s-new-in-infopath-2010.aspx.

Licensing

The SharePoint 2010 Licensing Q&A page can be found at sharepoint.microsoft.com/en-us/buy/Pages/Licensing-Details.aspx.

A comparison of the SharePoint product line's different tiers of functionality is posted at sharepoint.microsoft.com/en-us/buy/Pages/Editions-Comparison.aspx.

Microsoft's quarterly Product Use Rights document, which outlines rights and restrictions affecting volume licensing program customers, is available via a link on www.microsoft.com/licensing/about-licensing/product-licensing.aspx. The chapter "Microsoft Servers—Server/CAL" deals with SharePoint Server 2010 and FAST Search Server 2010 for SharePoint. The "Specialty Servers" chapter covers Search Server 2010 and both the Standard and Enterprise editions of SharePoint Server 2010 for Internet Sites.

Microsoft's "Product List for Volume Licensing" document, updated monthly and posted at www.microsoftvolumelicensing.com/userights/PL.aspx, explains some SharePoint-related licensing details such as downgrade rights and special accommodations made for Project Server customers who have added SA to Project licenses.

The advantages of using SQL Server Enterprise Edition with SharePoint Server 2010 are covered in the Mar. 2010 white paper "Microsoft SQL Server 2008 R2 Enterprise Edition and Microsoft SharePoint Server 2010: Better Together" available at technet.microsoft.com/library/cc990273.aspx.

A "Microsoft SharePoint Server 2010 for Internet Sites" volume licensing brief, dated June 2010, is available at download.microsoft.com/download/3/D/4/3D42BDC2-6725-4B29-B75A-A5B04179958B/Licensing_SharePoint_for_Internet_Sites.docx. The document provides five scenarios involving access by users external to an organization, along with a list of the required SharePoint licenses.

Using MSDN to license development or test environments involving SharePoint and other products is covered in "Understanding Licensing for Software Development" on page 21 of the Mar. 2010 *Update*.

Licensing details for SQL Server 2008 R2 are discussed in "SQL Server Gets More Editions, Higher Prices, Memory Caps" on page 24 of the June 2010 *Update*.

Licensing details for Windows Server 2008 R2 are covered in the Nov. 2009 *Licensing Outline*, "Windows Server 2008 R2 Packaging, Licensing, Pricing."