Activity-based Management of IT Service Delivery

John Bailey, Eser Kandogan, Eben Haber, Paul P. Maglio,
IBM Almaden Research Center
650 Harry Rd.
San Jose, CA 95120
+1 919 493 0628
{baileyj,eser,ehaber,pmaglio}@us.ibm.com

ABSTRACT
Growth, adaptability, innovation, and cost control are leading concerns of businesses, especially with respect to use of information technology (IT). Though standards such as the Information Technology Infrastructure Library (ITIL) offer the potential for cost savings through the use of formal processes and best practices, such top-down approaches tend to be either high-level – often far removed from the actual work – or low-level – often inflexible given the rapid pace of technology and market change. We conducted field studies to examine work practices in IT service delivery. Our results suggest that unstructured work activities comprise a significant and vital portion of the overall work done by people in IT service delivery. These activities include negotiating work items and schedules, seeking and providing information and expertise, and using and sharing custom tools and practices. Unstructured activities are conducted in parallel to formal, structured IT service processes, but are not well supported by existing integrated tooling. Thus, they are not easily accounted for and rarely result in reusable assets or feedback to improve the formal IT processes. Based on these findings, we propose an administrator workspace aimed specifically at blending structured and unstructured work activities to support effective, reusable, and quantifiable IT service delivery.

Categories and Subject Descriptors
H.5.2 [**User Interfaces**]; K.6.4 [**Systems Management**]

General Terms
Management, Human Factors, Services.

Keywords
Activity Management, IT Service Delivery, Workspaces.

1. INTRODUCTION
Today, IT services are delivered through a mix of structured and unstructured work activities. Structured activities rely primarily on standardized processes, procedures, and tools. In IT service support and delivery, an increasingly popular standardization effort is embodied by Information Technology Infrastructure Libraries (ITIL) [2, 3], which prescribe processes for capacity management, availability management, service-level management, and financial management to achieve high quality IT services. Unstructured activities involve local work practices, custom-developed tools, ad hoc collaborations with colleagues, and informal procedures. These activities include seeking information from colleagues or external sources, writing custom scripts, troubleshooting, renegotiating policies, seeking approvals, discovering change impacts, and sharing information on demand. Results from our studies of IT service delivery suggest that considerable time is spent on unstructured work activities [4, 6], which is not surprising as knowledge work is typically collaborative, informal, and situated [7].

As concerns over compliance, quality, and cost grow, IT companies often seek to standardize processes. Improving quality and efficiency of structured activities through standardization and automation has received much attention recently, simply because availability of workflow-tracking information has provided an opportunity to improve services. Yet, the complexity of information technology and the speed of change often outpace our ability to standardize optimal processes, making IT service delivery a highly customized activity. Simply put, top-down designed processes, implemented in workflow systems, cannot standardize all work activities [5]. As unstructured work activities are often not well understood, they are not accounted for in service delivery cost structures. Yet, without an accounting of unstructured activities, true cost of delivering service cannot be known. Focusing only on standardization without considering support for and integration with unstructured activities misses the opportunity for vast improvements in efficiency and quality of IT service delivery.

Current approaches in IT service delivery have many limitations, including:

- Top-down process design limits discovery and dissemination of local best practices.
- Rigid processes fail to provide sufficient flexibility to describe ad hoc steps and actions that are necessary in a complex, custom environment.
- It is difficult for practitioners to capture ad hoc or custom extensions to workflows.
- Reuse is limited, as only well-known, high-level processes are encoded, leaving out low-level practices, where real expertise is applied.

CHIMIT'07, March 30–31, 2007, Cambridge, MA, U.S.A.

- Capturing practices requires planning and deliberate efforts by the administrator and often occurs after the fact.
- Collaboration occurs outside current management environments; thus, ad hoc discussions are not preserved along with associated artifacts.
- Finding expertise on a particular problem is difficult, as tools do not capture and facilitate finding previous interactions with systems.
- Most system management tools are not integrated; thus, it is not possible to mix and match and organize low-level tasks by current activity.
- Integrated environments fail to preserve the business application context tools run in; thus each user initializes each tool over and over with context-related configuration information.
- There is no easy and accurate way to account for time and effort spent on unstructured activities.

In what follows, we first present two case studies drawn from our field studies of large IT service delivery organizations. Next, we introduce and describe in detail a concept design for Activity-based Management of IT Services (AMITS), which is an integrated workspace for IT service delivery personnel designed to address the above mentioned limitations. We conclude with a short discussion of the value prospects for the current and future AMITS designs.

2. CASE STUDIES

2.1 Introduction

Over the past several years, we have conducted field studies of system administrators in large IT service organizations in the US (see also [4]). Our method has been ethnographic, observing the minute-by-minute work of system administrators and interviewing system administrators in their normal work settings. Our analyses of these ethnographic data reveal how system administrators really work, including how they spend their time, who they interact with, and what tools and artifacts they use. Here, we provide two use cases that illustrate how personnel in IT service delivery organizations work with each other, using both specialized and common tools to complete structured and unstructured activities.

2.2 IT Outsourcing Case Study

We conducted a study of a team responsible for transition from initial IT service outsourcing engagement to a steady-state, operational system. We made two visits during the spring of 2005 to an enterprise IT outsourcing team located in a United States Midwest city. The first of these visits was as part of a larger investigation team that was conducting interviews of IT outsourcing personnel. The second visit was specifically for conducting observational research. In this second visit, we combined video-taped observations with interviews over a three-day period, resulting in over eight hours of video footage and several pages of typed notes.

2.2.1 Overview

The enterprise IT service outsourcing business we studied comprised teams responsible for different phases of the lifecycle of an outsourcing agreement. The phases include sales, engagement, transition, and production. The sales phase involves selling IT service to a client. The engagement phase involves working with the client to define details of a contract for services, including technical solutions, terms, conditions, and costs. The transition phase involves provisioning an IT solution, as defined in a solution architecture document. The production phase involves bringing the solution online to steady-state operation. Our study focused on personnel working in the transition and production phases. We observed and interviewed employees of the outsourcing provider in several different roles: transition manager, operations transition lead, operations analyst, operations team lead, and operations support analyst. The transition manager was the bridge between the engagement phase and the production phase. During the transition phase, she worked with the solution provisioning personnel to make sure the servers were built and configured to specification, and she engaged early with the operations team to ensure a smooth transition to production. The operations transition lead, operations analyst, operations team lead, and operations support analyst were all part of the production phase. Their primary responsibilities were ensuring that new servers “go live” successfully and continue to function properly once in production. As we will show, each interacted with one or more of their colleagues in context of a combination of structured and unstructured work activities.

2.2.2 Transition Manager

Xiang was a transition manager. We observed her on what she described as a typical day, as she was managing a large project that was partitioned into ten discrete phases. In the phase we observed, the teams were setting up or *provisioning* more than forty servers. Xiang used a large multi-tab spreadsheet to track all information about each server that was going to be provisioned. She created new versions of the spreadsheet manually and saved them in a central database. She received many updates throughout the day, and would save them until there were enough to justify a batch update, which resulted in a new version of the spreadsheet. The spreadsheet was large and complex, making it difficult for unfamiliar users to find information in it. Thus, Xiang responded to many phone calls, instant messages, and email messages inquiring about information contained in the spreadsheet.

Using instant messages, phone, and email, Xiang stayed in frequent contact with the lead architect and other team leads working on the solution specification, to request and receive updates for the spreadsheet, and to ultimately determine when it was complete. When Xiang had determined that the spreadsheet was complete, she opened a change ticket and assigned it to herself with a 35-day target. She cut and pasted the server configuration information from the spreadsheet into the change ticket. She then scheduled a kick-off meeting for those who would actually be provisioning the servers during transition and supporting the servers when they went into production. People who attended this meeting were from provisioning teams, operations, network support, and other functional areas. During the meeting, Xiang presented all information about the servers to these teams and answered questions.

After the provisioning work began, i.e. building out the servers, Xiang tracked progress on individual work items that were assigned to the different teams. There were teams to install the OS, create partitions, install middleware, and to handle networking. Because many of these are interdependent, such as configuration settings and order of tasks, Xiang helped to coordinate the work among the provisioning team members using the phone, email, and instant messages. Once the servers were built and the applications were loaded, Xiang worked with the operations transition lead to prepare operations for production. This included having a test period to make sure all the servers were working and that the operations team had access.

2.2.3 Operations Transition Lead

Arnold was an operations transition lead, responsible for preparing the operations team for the production "go-live" day for every new account. He worked closely with Xiang, beginning after the kick-off meeting, and accelerating the pace after the application management team finished installing applications. He also worked with the operations analysts on his team and with the operations support analyst. If there were special requirements for an account for hardware or software not already available in the operations center, then Arnold had to get that request in early to the operations support analyst.

Arnold was responsible for initial entry of account information, including names and contact information, server names, IP addresses, and physical location, monitoring policies, and service level agreement targets. The information needed to be entered into the correct databases, and the relevant operations team needed to be notified when it was available. The operations analysts were given their login and password information for the queues that they will support.

2.2.4 Operations Analyst

Rafey was an operations analyst. We observed him as he monitored production servers using two different tools, one for proactive monitoring of trends and another for alerts. He also read email and monitored his assigned problem ticket queues. Sometimes, problem tickets were opened by help desk analysts on the wrong problem ticket queue, and Rafey would have to transfer the ticket to the correct problem queue. If the queues were in the same version of the problem management system, it could be done electronically; otherwise, Rafey had to do it manually by cutting and pasting information into an email message, and sending to an analyst for the correct queue. The operations team had a standard goal of 90% problem resolution. The first thing that Rafey did with a problem ticket was to verify that the problem was still there. If it was, and if Rafey could not resolve the problem, then he escalated it to the appropriate level-two support person, such as an OS systems administrator, a database administrator, a network administrator.

If a service level agreement target was missed for a high severity problem, then at the end of his shift, Rafey manually copied information from the problem management tool to a special executive reporting tool to be included in a daily report. The project executives wanted to know about all missed targets as any one might become a critical customer situation.

2.2.5 Operations Team Lead

Theresa was an operations team lead, the first point of contact for her operations team. She was responsible for maintaining customer account documentation, such as contact information, updates to server information, and monitoring and response policies. She also maintained team operations data, such as shared monitoring IDs and schedules. She was developing a training program for on-boarding new operations analysts.

2.2.6 Operations Support Analyst

Gilad was an operations support analyst. He provided technical support for operations, including procurement of hardware, software, network, and telecom services, and installation of software and hardware. If there were special software requirements for account monitoring, then Gilad would procure and install it in operations. He coordinated his activities with the transition manager and operations transition lead.

2.3 Storage Management Case Study

In 2005, we made two visits to study storage administrators at a large government facility that maintained several petabytes of data, with storage growing by several terabytes per day. The sheer volume of data means that most files are stored on tape most of the time, and are moved onto disk only when needed. For frequently accessed data, this process is automated (for the most part) by several Hierarchical Storage Managers (HSMs), which include huge robotic tape silos and large disk arrays, with automated controls to move data between the two as needed. In the best case, it appears to the user that data is always available on (a sometimes slow) disk, though tape robot jams and other system failures occurred almost daily, requiring systems administrator intervention. We primarily interviewed and observed storage system administrators that managed HSMs. We were not permitted to videotape their work.

2.3.1 Storage Systems Administrators

We observed Mandy and Arthur, two storage system administrators. They worked very closely together as part of a four-person team managing one of the HSMs, and were often in the same office discussing and diagnosing problems.

Mandy's work day involved handling requests that came in via the online trouble ticket system (e.g., "I can't access file ZZZ"), ensuring that new machines have access to the appropriate data, moving data between HSMs when necessary, and cleaning up files when staff members leave the organization (either moving the data to long term storage, or making it available to a successor staff member). Mandy also collected and analyzed performance data to make sure her HSM was working as expected.

Configuring new machines, data movement, and clean-up were fairly structured tasks. Trouble tickets were idiosyncratic, however, requiring considerable investigation and judgment to determine the cause and appropriate solution. Problems might be caused by users pulling in too much data (overflowing the disk cache), software configuration errors, hardware failures, or even dirty or peeling labels on magnetic tapes (found by a trip inside the robotic tape silo). In the case of failures caused by vendor-supplied hardware and software, Mandy was responsible for contacting and working with vendor Service Engineers (SEs), including filling out the necessary paperwork to give the SE access to the appropriate buildings and rooms. Problems might require minutes, hours, or even days to resolve. We observed one particularly intense troubleshooting session in which Mandy and Arthur tried to determine why an entire tape silo was crashing. This involved collecting log and configuration data from the silo and other interacting systems, writing test programs to simulate various read/write loads, speaking with the silo vendor's customer support, paging and getting building access for the SE, and experimenting with different configuration values. The problem

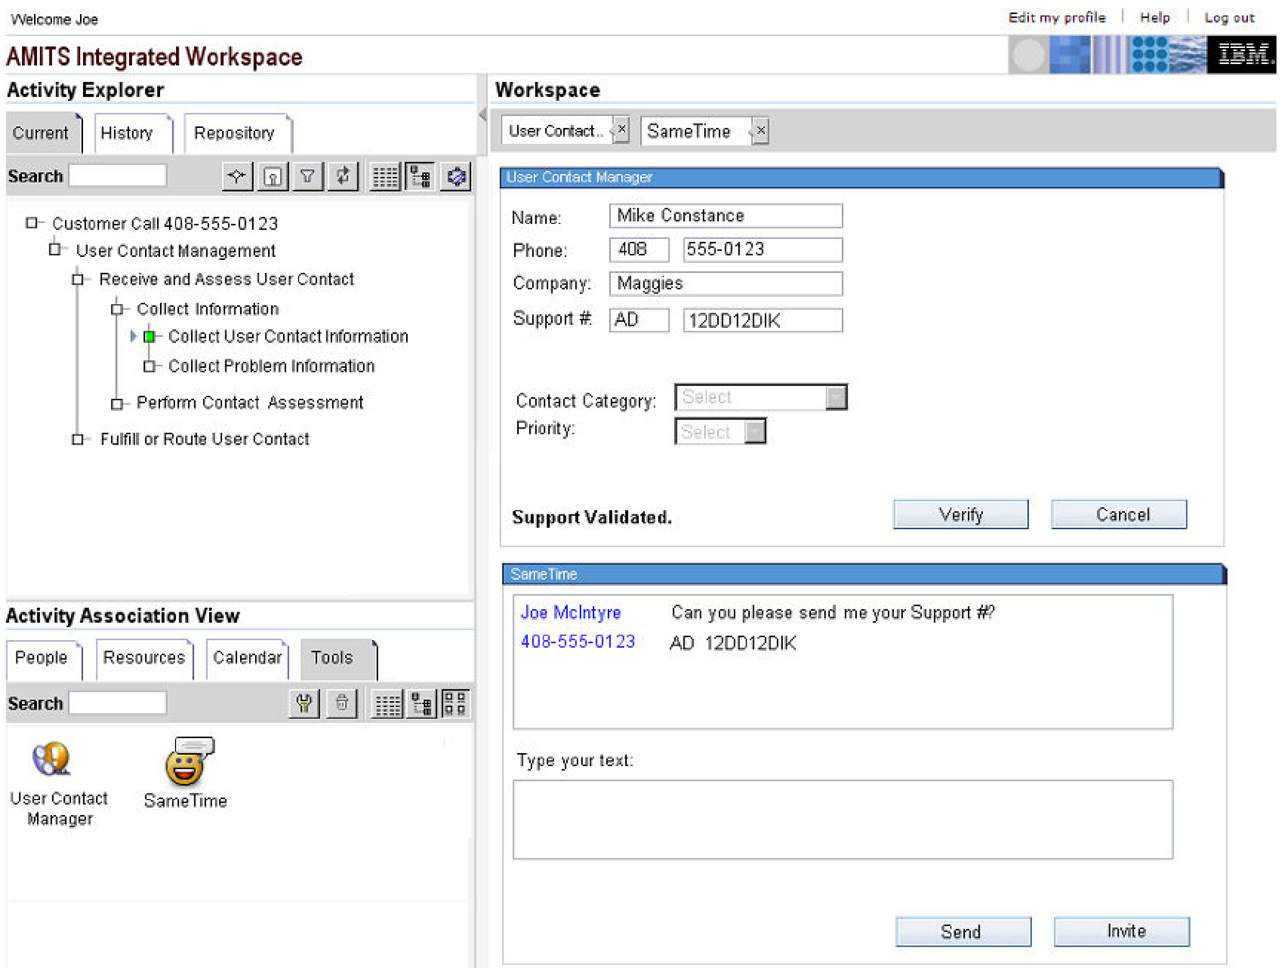

Figure 1. AMITS integrates process, people, time, and systems management in Activity Explorer, Activity Association, and Workspace views, respectively.

had many causes, including a defective robot hand, an incorrect configuration parameter, and an end-user application error.

Arthur's responsibilities were similar to Mandy's, but he had somewhat more experience and spent more time building tools for monitoring and performance analysis. To make their work more proactive, Arthur created an entire infrastructure of automatically running tools that continually checked for failures or performance problems, e-mailing or paging him when something went wrong. Some failures could be handled automatically by his tools, others required human intervention. Arthur had a test system on which he could introduce failures to see how his tools reacted. Arthur's tools worked well, though there were no procedures in place to share them with other groups so they could be used more widely.

2.4 Analysis

The first case is quite broad, involving a large number of people working in different capacities. The second case is rather narrow, involving two people working in a small, specialized team. However, they are alike in that both describe settings in which people create and share information, and coordinate joint work. In both cases, structured work depends on standardized tools and processes, such as those for creating and managing change tickets. And in both cases, people coordinate joint action through communication that goes *around* these standard processes, such as through instant messages and email to check on server settings. Sometimes, people rely on tool-supported processes to structure and create auditable records of their work, and sometimes, they use the phone, email, and instant messaging to collaborate. But often, information created during collaboration – or even the record of the event itself – is lost.

For example, one of the steps that Xiang took during transition was to create a change ticket and assign it to herself. The change ticket contained descriptions of all the servers being provisioned, but because this description was created outside of a formal audit trail for a block of work, it was of limited use. Xiang was responsible for negotiating completion dates from each team, tracking progress, assisting in coordination, and finding and resolving problems that were impeding overall progress. She did this by proactively using phone, email, and instant messages. If a problem occurred, Xiang might not have found out about it for several hours or days. She had very little visibility into what was happening and what had happened. Other provisioning teams also had very little visibility onto what each team had done, any problems encountered, and what their next step would be.

3. DESIGN

The AMITS approach to integrating structured and unstructured work activities is based on the Unified Activity Management (UAM) project [1]. Because most work, whether individual or collaborative, cuts across many independent tools, UAM proposes an explicit representation of work to support users to setup, perform, coordinate, and plan work activities effectively. In UAM, activities are abstract computational constructs with properties and operations.

AMITS relies on key ideas from UAM, applying and extending them for the IT service delivery domain. Central to our approach is a workspace that features activities as first class objects, seamlessly integrating processes, people, schedules, and resources (Figure 1). In the context of IT service delivery, processes refer to the structured activities supported in formal workflow systems, such as problem management and change management. In AMITS, they are represented in default activity instance structures, as generated from process templates. People are defined by roles and individuals that participate in the structured and unstructured aspects of an activity, with effective collaboration – seeking information from a colleague or customer, finding expertise, establishing common ground via shared context, and coordinating tasks – being a primary focus. Schedules, represented in AMITS via calendar views and activity status, are particularly critical in IT service delivery because the effective use of time and resources determine whether contractual obligations and cost targets are achieved. Resources, generically describing anything used by or for an activity, are presented in AMITS as tools and resources, with resources more narrowly defined as documents, links, or any other type of information artifact.

3.1 Activity Representation and Instantiation

New activity instances can be created from structured templates, or they can be created as ad hoc activities. For example, on receiving a call from a customer about a problem, a service desk operator can create an instance of a standard ITIL User Contact Management process. The formal ITIL process steps, along with all the associated resources, roles, and tools are imported into the workspace. In another example, a capacity planner might research the impact of new disk systems by jotting down a few tasks and then expanding the list with more detailed steps over time. Activities can be instantiated manually or automatically in response to system events and scheduled work. For example, a troubleshooting activity can be automatically instantiated on a server-down event and a backup activity can be instantiated regularly each week.

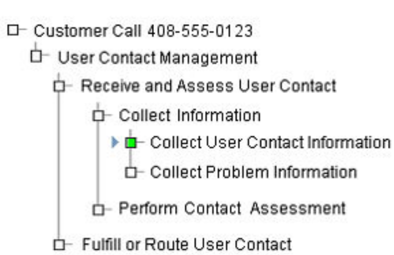

Figure 2. Upon receiving a customer call a new activity is instantiated from the ITIL User Contact Management.

When a new activity is created, the user's workspace shows the activity in the Activity Explorer view (Figure 1). Though there are many possible representations of an activity, including graphical and tabular representations, the preferred representation is a simple hierarchical listing of activities (Figure 2). In this representation, activities are rendered as a tree, with each node of the tree containing the name of the activity. Users can browse the activities by expanding and collapsing nodes. Each node in the tree, except leaf nodes, is expandable and collapsible to show and hide subactivities of the selected node.

When a node is selected, it can be set as the current activity. In this case, a triangular shape mark is put next to the expand/collapse box. As a result, the Activity Association views and tools panel are updated to show the people, resources, tools, and calendar items of the current activity. The content of the views are updated as users select different activities in the hierarchy, thereby creating a work context. In the default case, only items that are directly associated with the current activity (e.g. people, resources) are shown. Alternatively, users can choose to show recursively all associated items of the selected current activity and its subactivities.

In the People view, selectable by clicking on the People tab in the Activity Association views, users can see people associated with the activity, along with their roles. In the default view, associated people are shown in a hierarchical organization of roles. Alternative renderings are possible, such as iconic representation of people with detail attribute notations, such as instant message status, participant project responsibility, expertise, and online availability. Various actions can also be connected to entries in the People view, such as open instant messaging with selected participant, view participant activity information, and assign tasks to participants.

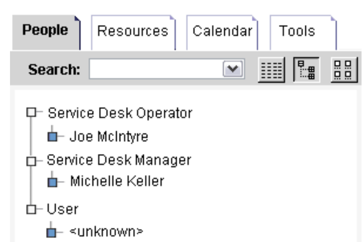

Figure 3. Associated people of an activity are shown hierarchically by their role.

The Resources view likewise shows all the resources associated with the current activity. The default representation of the resources is a table with various columns of resource properties, such as resource name, creation-modification time, length, and size. Alternative representations, such as a hierarchical view by resource type and activity type can also be useful. Various actions can be connected to resource entries, such as view/edit resource, (un)lock resource, and view edit history.

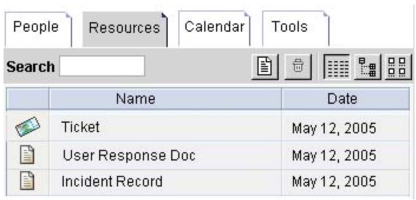

Figure 4. Resources associated with an activity.

The Calendar view serves two purposes: to view past actions associated with an activity, and to view and schedule future scheduled events, to-dos, and deadlines. The basic representation is a tabular representation of actions along with start date and time and duration of the actions. Alternative representations, such as daily, weekly, and monthly calendar-like representations and time-line representations could facilitate easier scheduling and planning. Scheduling can be as simple as reminders and as complex as monitors that can trigger activity instantiations. For example, a change management activity that is not complete by the due date can be automatically escalated, notifying both the activity owner and the change manager. In another example, a recurring calendar to check a system health could automatically instantiate an instance of an activity. These kinds of functions would have been very helpful to Xiang, the transition manager in the earlier case study, in tracking progress of provisioning activities and allowing her to manage on an exception basis. Moreover, the operations transitions lead could be a participant in these activities, monitoring progress and getting an early start in preparing for the "go live" date.

In the Tools view, tools associated with the current activity are shown. Various representations of tools are possible, including iconic and tabular listings, and perhaps tool usage attributes such as frequency by task and last use. Tools launched from the Tools view are displayed in the tools workspace.

Each activity can be set with various status indicators that describe the urgency, priority, severity, and completeness. Though some of these status indicators can be set manually, it is also possible to automatically update activity status. One way to do this is to automatically escalate an activity based on a scheduled event. For example, a scheduled backup activity can increase the priority of a system configuration activity that is required to precede the backup work.

Once users have finished their participation in an activity, they can mark it complete. Consequently, the status of the activity is recorded and synchronized to all other participants of the activity.

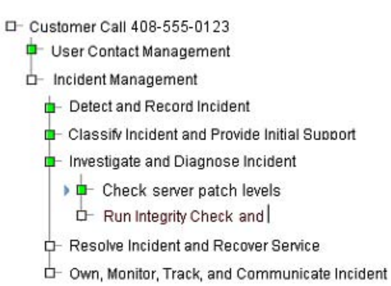

Figure 5. Users can create a new subactivity at any point by inserting a new node into the tree with activity title.

Note that current activities are automatically saved to the repository both when users work on the activity and when it is completed, facilitating synchronization with activity participants. Saved activities can also be searched for and reused in new activities by anybody who shares the repository.

3.2 Activity Capture and Extension

Once an activity is instantiated or created, users often see the need to extend the activity by adding new subactivities to detail specific tasks, add or modify associated people based on the current project team, update documents and schedules, and add their own favorite tools. This can be done manually by selecting specific subactivities, people, resources, or tools, automatically through various algorithms to match tasks, skills, prior experience, or tools, or organically through interaction with the tools, people, and resources within the context of the activity.

Manual extension of an activity is potentially useful when an activity is initially instantiated from a template activity in the repository. In this initial planning phase users can put in schedules for known activities, team members, links to artifacts, and required tools. Often users provide known team members by filling in specific roles in the template activity. Also at this stage, people may find it necessary to extend the activity for the current task by adding further tasks to do, customizing existing activities, and providing a schedule for various activities. Each activity can be associated with people, tools, resources, and schedules. After the initial phase, it is expected that most of the extensions will take place as the activity organically grows through interaction within the context of the current activity, as users use tools, chat with people, create artifacts, or otherwise engage in work. All these tools, people, and resources will be associated with the current activity. Once an activity has a sufficient amount of content, automatic extensions to activities are possible through various matching algorithms that can be executed on demand by the user to find and import activities with similar tasks, people who might provide expertise, and tools that were used effectively in similar activities.

The above AMITS capabilities could considerably enhance the effectiveness of Mandy, the storage systems administrator introduced in the second case study. For example, as part of Mandy's troubleshooting activities, some types of problems required contacting vendor service engineers. Based on matching algorithms of prior problems, these instances could be quickly isolated and related vendor documents containing contact information, support codes, and availability could be imported into the current activity.

Adding a new subactivity to an activity is accomplished by selecting an existing activity and inserting it (Figure 5). Users would then provide a name for the subactivity, and potentially associate people, resources, and tools. Likewise, subactivities can be deleted and modified at any point in the activity tree.

To add a person manually, one can either search for a known person or browse the specialist repository within an organization (by name, project, role and/or expertise) and associate to an activity or subactivity (Figure 6). There are also various ways to automatically add a person to an activity. When an activity is instantiated from the template activity repository, people and roles associated with the activity are automatically associated with the new activity as well. Potentially, a role that is associated with a subactivity that is imported could also be matched during instantiation. For example, when a service desk operator, instantiates a new activity from the User Contact Activity template, automatically a Service Desk Manager is associated with the activity and matched to the user's manager. Collaboration and communication actions in the context of an activity can also automatically associate the parties involved to the current activity. For example, when a specialist consults an expert during the course of an activity, the expert is automatically associated with the activity. In other cases, the expert might not be known at the time but searched through the skills and activity matching utility. In this case, the expert is found based on experience with similar activities by crawling the activity repository and performing a similarity matching.

In a similar fashion, resources are created, either organically through the use of the tools, automatically from the repository (from templates or by activity type), or manually created from resource template repository. Once created, they are added to the Resources view with the appropriate time stamp. New resources can be added from the resource repository through either search or browser interface. The resource template repository can also be browsed by resource name, type, and by resources organized by activity type. The search interface allows users to find resources conveniently from the resource template repository. Likewise, tools can be associated with an activity either manually by launching the tool in the workspace from the tool repository, or automatically from activity templates or by activity type (Figure 7).

Figure 6. During interaction resources and tools used are captured and associated with the current activity.

As users interact in the workspace, instantiating a new activity from templates, providing detailed steps to perform a system management task, adding (assigning) people to activities, looking up resources, and interacting with system management tools, each action is captured and saved along with the activity in the Calendar view (Figure 8). As a result an activity spawned from a template-structured process and extended by user actions taken in context of the structured process are captured and available for reuse when searched in the activity repository, which contains all activities performed in an IT service delivery center.

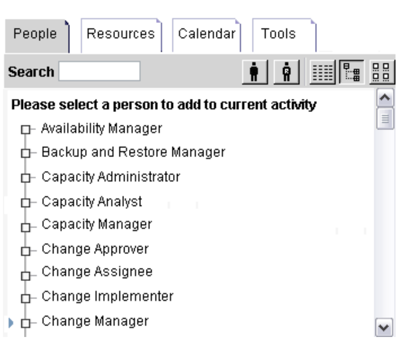

Figure 7. New people can also be manually associated to an activity.

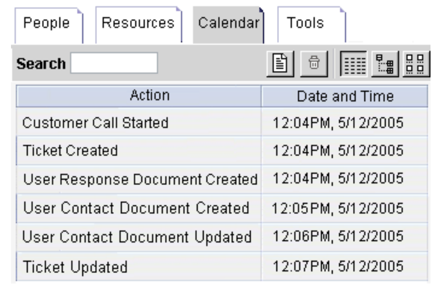

Figure 8. All actions are recorded in the Calendar view to provide an accurate picture of activity history.

New activities created from scratch, and used over time in different circumstances, can selectively evolve into standards, becoming part of the activity templates used by the larger community in an organization. As users interact in the workspace – performing the detailed steps of IT service work, assigning and collaborating with people (e.g. through IM, e-mail, or IP phone), creating and working with resources (e.g. documents, spreadsheets, web pages), and interacting with system management tools (e.g. through user actions) – each person, resource, and action is captured and associated with the current activity, providing a complete picture of the work conducted. Each of these actions are captured into the repository and also recorded on the calendar with respective times of start and completion. As a result all current and past activities in an IT service delivery center are stored in repositories, which can be searched to facilitate reuse by incorporating matching activities into new work.

The Activity Repository contains all current and past activities within an organization (Figure 9). There can be multiple activity repositories based on location, customer, delivery center, competency, and they can be federated, or brought into a specialized service offering.

The main purpose of the repository is to have a single location to capture organizational knowledge within a group, so that practices can be shared and reused. By default the activities in the repository are organized by activity type, in a hierarchical fashion but other representations are certainly feasible, such as organized by activity name, date, project, or person/role.

A person's past activities can be found in the activity history. The Activity History view is a convenient way to browse and search past activities for one's own activities. It can be browsed and searched by keyword, name, persons involved, dates of the past activities, and other useful attributes. Using the activities found in the Activity History, specialists can refer to prior actions, examine past solutions, and contact prior collaborators. If desired subactivities from the History View can be imported or inserted into the current activity.

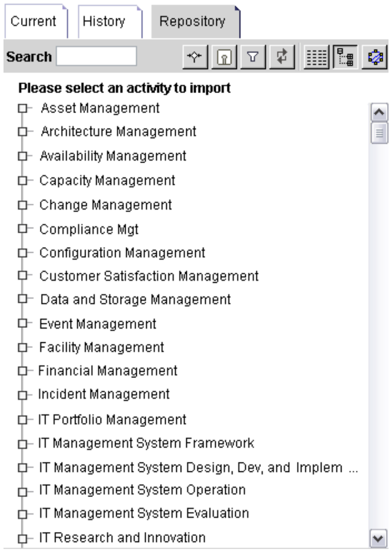

Figure 9. Activity Repository is a store for all standard and ad hoc processes captured.

3.3 Tools Workspace and Contextualization

When users launch tools from a Tools view, the launched tool is displayed in the tools workspace. Tools workspace is like a set of tool windows. Tool windows can be arranged, minimized, and resized, much as in all window managers. One notable difference is that only those tool windows associated with the current activity are displayed. It is more like a multi-desktop window manager, with each desktop corresponding to an activity. Though by default, only tools that are associated with the current activity are displayed on the tool panel, users can set it in such a way that tools belonging to subactivities of the current activity are also displayed at the same time.

Each tool in the tools workspace is launched in the context of the current activity. This means that tools can automatically set data in certain UI elements from the activity context resource. An activity context resource is essentially an XML document associated with every subactivity with the purpose of storing contextual information (Figure 10). Tools can read and write to the context resource by interaction – transparent to the users. Users can also manually edit the context resource just like editing any other document resource. For example, in the user contact management activity, when the call is received from the customer, an activity context resource is created automatically, containing the customer phone number. When tools are launched in the context of this activity with UI elements addressed to customer phone number they can read in this value (i.e. customer phone number) from the activity context resource and automatically fill in corresponding element. Tool developers (as well as specialists) can associate UI elements in the tool to refer to any data path in the activity context resource. Thus, if the activity context contains data referred by a path of the UI element in a tool, its value is automatically read from/written to the activity context resource. If the value is lacking and the user manually enters the value in a UI element, the user entered value is synchronized back to the activity context resource.

3.4 Activity Search and Reuse

The Activity Repository can also be searched in a number of ways. Keyword-based search is available for finding activities based on matching a keyword with content on the activity structure, resources, people, and tools. Alternatively, repository search can use a similarity-based algorithm. In this case, given a current activity, the repository search will return matches based on similarity of activity structure, content, people involved, resources and tools used. This could come handy as specifying a search query can be difficult at times, particularly in complex IT tasks.

Given these alternative ways to browse and search the repository, users can find a relevant past activity and perhaps refer to prior actions performed within the activity, examine solutions, and find collaborators. If desired, a past activity found in the repository can be used as a template for instantiating a new activity or parts of it can be imported into a current activity at any point in the activity structure. Along with the activity structure, all associated people, tools, and resources can be imported if desired. This could help the specialists a great deal when working on a new project, as specific past experience can be brought to bear using tools known to have worked in the past, using resources that helped colleagues, and providing access to colleagues who had actually done similar work (Figure 11).

In addition, activities can also be received by email or found in web documents. In these cases, activities in these documents are in the form of activity descriptors, which define activity structure and associated elements. Simply copying and pasting them into the current activity in the workplace automatically copies the copied activity structure.

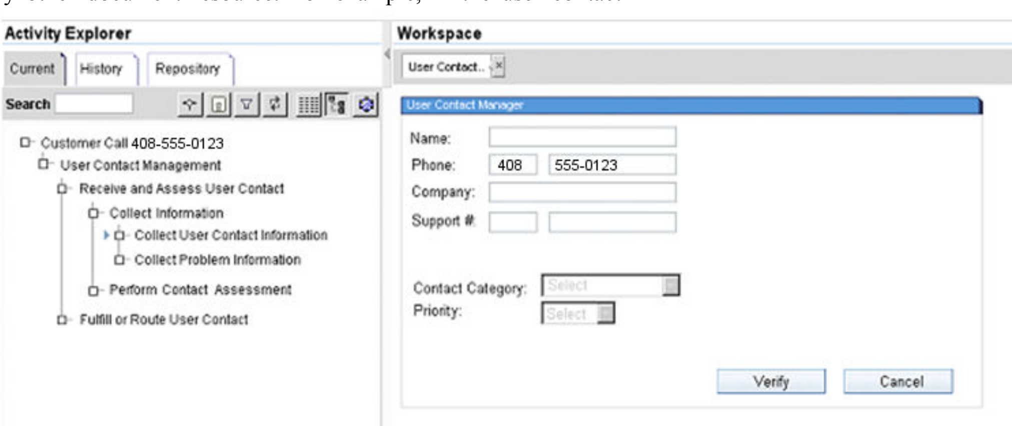

Figure 10. Workspace tools can inherit context from current activity. For example, user contact manager tool automatically sets the phone number to the customer phone number associated with the user contact management activity.

3.5 Activity Monitors and Triggers

Activities can also be automatically instantiated in response to system events (e.g. User Contact Management activity in response to support phone call). Basically, users can create custom system event monitors and associate a default activity from the activity template repository. Once event monitors are active and the pattern is detected, the default activity associated with the event is automatically instantiated for the registered user.

Another way to instantiate an activity is to schedule events in the Calendar View. Basically, an event is created for a future date and time and registered with a default activity from the activity template repository. Similar to system event monitors, when the time comes, the associated activity is automatically instantiated.

Users can also create monitors over the activity repository to track activities by type, urgency, priority, status, person name, and custom pattern. A monitor tracks updates to the repository and dynamically triggers an alert when monitor condition is satisfied. For example, a manager could build a monitor that tracks all change management activities and be notified when they are completed. While an alert can be as simple as an email or pager notification it could also automatically instantiate a new activity.

3.6 Activity Filters and Reorganization

Often activities are shared by multiple people. Based on the roles of the participants in an activity, different users may want to view the activity from different perspectives. For example, a database administrator may want focus just on the data management parts of a larger activity to setup an ecommerce site for a large bank. A delivery manager might be interested in all activities within the project but only higher level tasks within the larger activity. Depending on their roles, some users might want to focus on a number of activities but from a particular angle. For example, a change manager might be interested in all ecommerce site development activities but only as far as change requests are concerned.

To support users to create their own perspectives on activities, users can transform current ongoing activities. Basic transformation operations include select, create, copy, and aggregate. To create an activity that is a transformation, users begin by creating a new activity normally. At any point in the new activity structure users can create a new subactivity that is in fact an activity path expression (similar to XPATH) to the current activity repository. The path expression is built by demonstration where users would click on one or more reference activities in the repository and a path is built. For example, if the user has clicked

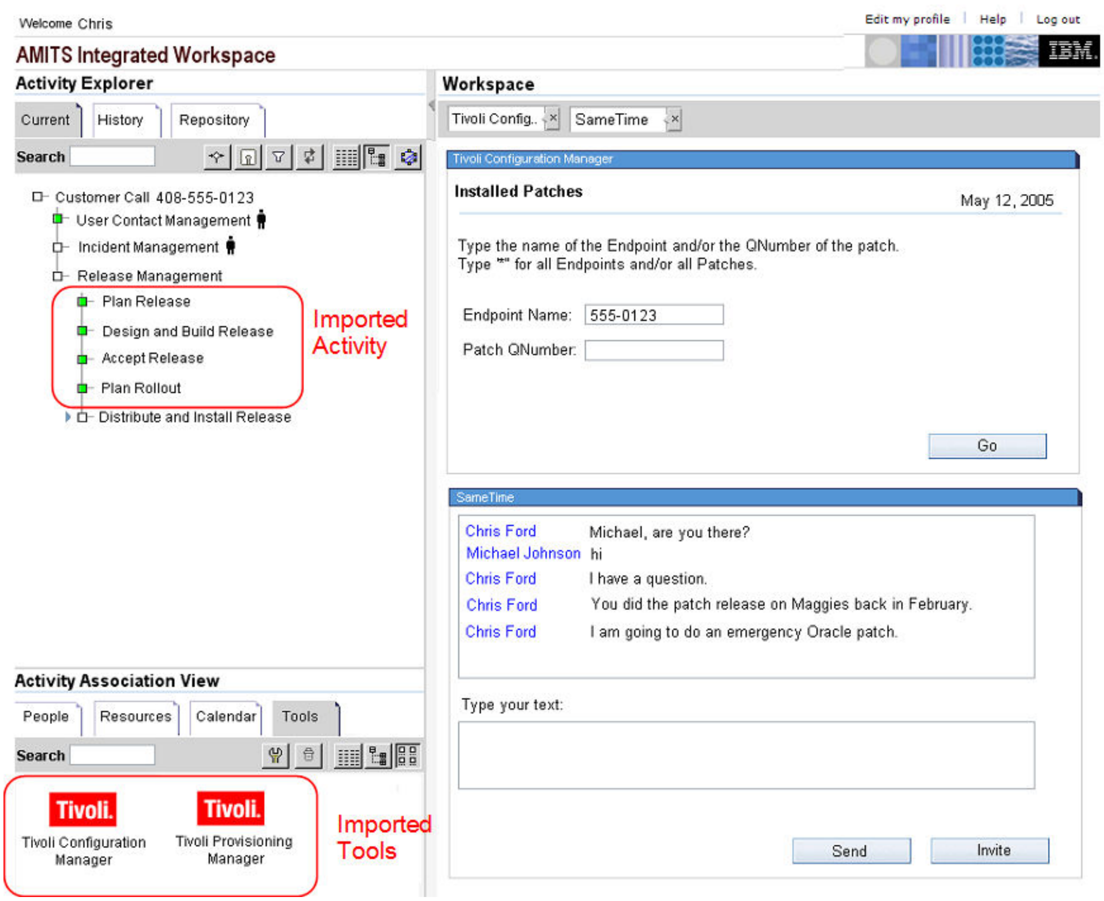

Figure 11. When an activity is imported from the repository, associated task steps, people, resources, and tools are inserted into the current activity workspace. This way users can substantially benefit from past practices easily.

on a number of user contact activities occurring at various points of the current activities, the activity path expression created would be something like `/*/user_contact`. Users might want to get only those user contact activities belonging to a particular project (defined in the activity context document), then the path expression would be like `/*/user_contact[project="abc"]`. If the user is interested in all subactivities of the user contact as well, the path expression would be `/*/user_contact/*`. Attributes of an activity can be set similarly through path expressions that also perform computation. For example, to count the user contact activities belonging to a project abc, one would write `count(/*/user_contact[project="abc"])` in the activity context document for a particular variable.

3.7 Activity Analysis

As users interact within the context of the current activity, tool actions, people interactions, and resources created and modified are also recorded in the backend for activity analysis, such as efficiency, quality, and cost analyses. Having the complete list of all actions carried out in an activity – rather than actions related to structured activities only – one can analyze most aspects of service delivery with much greater accuracy. Such analysis can be performed to build models of work, which can be used in cost analysis and planning, examine where time is spent, how much time particular type of activity takes and what is involved in carrying it out, and where things can be improved in general.

An enhanced monitor type is a report that creates a document with content and charts, driven by data from the activity repository. Reports are dynamically updated in response to monitor triggers. Reports can utilize statistical packages for analyzing cost, efficiency, and quality. In this case, data from the repository is processed by these statistical packages to arrive at aggregate metrics which are then rendered in the report documents.

4. CONCLUSIONS

The primary advantage of our activity-based service delivery approach is the coupling of structured and unstructured work, including workflow in ITIL with related informal activities and artifacts, which can collectively create a shared awareness of priorities, urgencies, and team performance in relation to common objectives. Availability of precise data about delivery activities could potentially lead to improved efficiency and quality of service delivery by providing better accountability for the actual cost of delivery.

While we believe that the AMITS concept designs hold sufficient promise to warrant further investment, we do not want to leave the impression that everything is solved. For example, potential users may resist using a workspace that records practically all of their interactions, while storing them into a shared repository. Also, while extending standard templates to include new steps that have been included frequently as unstructured extensions is technically simple, developing an acceptable governance model for this approach may not be so straight forward. Finally, some of the advanced analysis and matching algorithms that we've proposed pose significant challenges.

AMITS remains to be validated in practice, and has challenges to address, but integrating structured and unstructured activities offers substantial promise, leveraging basic human creativity, facilitating effective collaborative work and reuse, and offering more accurate planning and delivery of high quality, effective services. Future extensions to AMITS could include the integration of human activities with automated activities, with the activity-based approach serving as a framework for progressive authorization and delegation of work to automation systems.

5. ACKNOWLEDGMENTS

We thank Chris Campbell for help with the IT outsourcing study and Cheryl Kieliszewski for help with the storage management study.

6. REFERENCES

[1] Moran, T. P., Cozzi, A., and Farrell, S. P. 2005. Unified activity management: supporting people in e-business. Communications of the ACM 48, 12 (Dec. 2005), 67-70.

[2] Office of Government Commerce, UK. (2001) Best Practice for Service Delivery – ITIL The Key to Managing IT Services, Stationary Office Books.

[3] Hochstein, A., Zarnekow, R., Brenner, W., ITIL as Common Practice Reference Model for IT Service Management: Formal Assessment and Implications for Practice. IEEE International Conference on e-Technology, e-Commerce, and e-Service (EEE '05), 2005, pp. 704-710.

[4] Barrett, R., Kandogan, E., Maglio, P. P., Haber, E. M., Takayama, L. A., and Prabaker, M. 2004. Field studies of computer system administrators: analysis of system management tools and practices. In Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (Chicago, Illinois, USA, November 06 - 10, 2004). CSCW '04. ACM Press, New York, NY, 388-395

[5] Abbott, K., and Sarin, S. Experiences with workflow management: Issues for the next generation. In Proceedings of CSCW 94, ACM Press,1994, 113–120.

[6] Maglio, P.P., Kandogan, E., Bailey, J., Case Studies in IT Management: On Formal Processes and Informal Activities in Service Delivery. 15th Annual AMA Frontiers in Service Conference

[7] Suchman, L. Office procedures as practical action: Models of work and system design. ACM Transactions on Information Systems 1 (1983), 320–328.