Modeling and Analysis of Information Technology Change and Access Controls in the Business Context

Andrew P. Moore (CERT Program, Software Engineering Institute)
Rohit S. Antao (CyLab at Carnegie Mellon University)

March 2007

TECHNICAL NOTE
CMU/SEI-2006-TN-040

Survivable Enterprise Management Initiative

Table of Contents

List of Figures

List of Tables

Acknowledgements

Many thanks to our sponsors of this project at the Information Technology Process Institute—especially, Gene Kim and Kurt Milne—and to Julia Allen of the SEI for support and insights. The Survivable Enterprise Management team in the CERT Coordination Center® continues to be a great place to work and very supportive of new efforts to bring system dynamics modeling and analysis to bear on important problems in information system security, survivability, and resiliency. It has been a true pleasure and inspiration to work with such talented individuals. Finally, thanks to the paper reviewers from the 2006 International Conference of the System Dynamics Society for their comments and pointers to some very useful resources and related work. A version of this paper was published in the proceedings for that conference.

Abstract

Ongoing field work centered at the Information Technology Process Institute (ITPI) makes clear that processes that control change and access within information technology (IT) management and operations simultaneously reduce security risk and increase efficiency and effectiveness. The CERT® Coordination Center is building on this work. This technical note describes a system dynamics model that embodies CERT's current hypothesis of why and how these controls reduce the problematic behavior of the low-performing IT operation. CERT has also started to extend the model in ways that reflect the improved performance seen by high performers. In the longer term, the hope is that this model will help to specify, explain, and justify a prescriptive process for integrating change and access controls into organizations' business processes in a way that most effectively reduces security risk and increases IT operational effectiveness and efficiency.

1 Introduction

As information technology (IT) makes a large and more noticeable contribution to business success, senior executives are under mounting pressure to clearly demonstrate the business value of IT, and to prove that IT investments can generate a positive return while supporting business objectives [Sarvanan 2000, ITPI 2005]. In order to meet these objectives, they must identify and recommend a set of processes and controls that improve IT management performance. Our research at CERT® builds on foundational work done at the Information Technology Process Institute (ITPI) [Behr 2005]. The ITPI is an independent research organization that supports IT audit, security, and operations professionals [ITPI 2007].

This section describes what it means to be a high-performing organization, the foundational work done to determine the cause of high performance, and the goal of our current work.

1.1 IT RESPONSIBILITIES AND PERFORMANCE

With IT occupying an integral position in the operations of any modern business, it faces the daunting challenge of succeeding in an increasingly competitive marketplace and complying with stringent regulatory requirements [Castner 2005]. The IT department, being a business enabler in most modern organizations, is entrusted with two broad responsibilities [Taylor 2005]:

1. Operate and maintain existing services and commitments.

2. Deliver new products and/or services to help businesses achieve their goals.

In the process of fulfilling these responsibilities, IT operations are simultaneously presented with various demands. The need to ensure that IT aligns with business objectives has made it necessary for IT operations to not only get the job done, but get it done in an effective and efficient manner.[1] In addition to coping with demands of effectiveness and efficiency from businesses, IT departments must satisfy regulatory requirements issued by laws such as the United States Sarbanes Oxley Act of 2002 [SoftLanding 2005]. Such requirements mandate the presence of a strong internal control structure to manage any risks that IT poses.

IT performance indicators measure how well an organization's IT department is doing in terms of achieving the desired results. Based on the responsibilities as-

---

[1] *IT effectiveness* is the extent to which IT processes produce the desired objectives. *IT efficiency* is the extent of IT resources used and needed to achieve those objectives (Brenner 2002).

signed to the IT division and the demands placed on the way they are fulfilled, the Software Engineering Institute (SEI) and the ITPI have developed a set of high performance indicators, which are listed in Table 1.[2]

1.2 FOUNDATIONAL WORK

For over five years, researchers at ITPI have been studying high-performing organizations in order to understand their IT processes and implementations. They continue to observe that these organizations evolve a system of process improvement as a natural consequence of their business demands and address security in the normal course of operational business. Surprisingly, these high-performing IT organizations have independently developed virtually the exact same processes to better manage their operational environment in order to achieve the desired performance outcome [Behr 2004].

Table 1:   High-Performance Indicators

Deliver new projects		Operate / maintain existing IT assets	
Effectiveness		**Effectiveness**	
1	High perceived value from the business	1	High uptime and service levels
2	High completion rate of projects, on time and on budget	2	Satisfactory and sustained security
3	High customer/user satisfaction with security	3	Low amounts of unplanned work
		4	High change rates
		5	High change success rates
		6	Low number of repeat audit findings
1	High application developer to completed project ratio	1	High server / system administrator ratio
2	Low % of development cost on security	2	High first fix rate
		3	Low % of IT budget spent on compliance
		4	Low % of IT budget spent on operations

More recently, ITPI began working with the SEI to better understand how these organizations manage IT to achieve business objectives, and to identify the core set of controls they rely on. *Controls* are processes that provide assurance for information and information services, and help mitigate risks associated with technology

---

[2]   Allen, J.; Behr, K.; Kim, G. et al. *Best in Class Security and Operations Round Table Report.* Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2004. Not publicly available.

use. Based on experience, ITPI hypothesizes that not all IT controls contribute equally to IT effectiveness, efficiency, and security. Those IT controls that contribute most significantly we call *foundational controls*; they help address operational effectiveness, efficiency, and security simultaneously.

In order to test this hypothesis and identify the set of foundational controls, ITPI launched the ITPI IT Controls Benchmarking Survey, which inquired about organizations' use of IT management controls, including change controls and access controls [ITPI 2004]. *Change controls* are controls that ensure the accuracy, integrity, authorization, and documentation of all changes made to computer and network systems. *Access controls* are controls that ensure access to systems, data files, and programs is limited to authorized users (IIA 2004). The survey spanned 89 organizations as of October 2005. Preliminary results of ITPI's analysis of data from this survey indicate a strong correlation between change and access controls and the high performance seen by some organizations. It also shows change and access controls to be foundational.

The field work indicates that high-performing organizations view change and access controls as critical to organizational success [Behr 2004]. High performers believe that these controls not only help satisfy regulatory requirements, but actually facilitate achieving the performance levels they desire. While these findings are encouraging, researchers observe that low-performing organizations also implement change and access controls, but they argue that these controls are useful primarily in satisfying regulatory requirements. When faced with performance problems, the low performers believe that change and access controls only serve to hinder recovery and must be circumvented to get work done faster. Ultimately, they believe that these controls are overly bureaucratic and diminish productivity [Kim 2005].

1.3 OUR RESEARCH

Motivated by the conflicting positions on the efficacy of change and access controls in IT performance, we attempt to determine causal relationships between change and access controls and IT performance. We hypothesize that a root cause for the performance problems experienced by many organizations lies in a tendency to relax the enforcement of change and access controls and shift excessive resources from proactive to reactive work to deal with system disruptions.[3] This behavior arises from an inability, or even negligent failure, to take into consideration the long-term effects or unanticipated consequences of the decision to bypass these

---

[3] Moore, A. P. & Antao, R. *System Dynamics Modeling and Analysis of IT Management Controls in Context*. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2005. Special Report. Not publicly available.

controls. There is a disproportionate focus on short-term profits as opposed to long-term improvements.

An uncommitted patchwork approach to the implementation of these controls makes them ineffective, thus preventing organizations from deriving their true value. This inevitably results in these controls being viewed as unnecessary over-head and, therefore, detrimental to productivity. This work attempts to provide a holistic view of the IT operational environment with respect to change and access control management. Armed with this enhanced understanding we develop an appreciation for the improved operational performance that these controls can bring about. Table 2 indicates some of the benefits we hope to achieve through this work.

With an improved understanding of how these controls can be used to make everyday operations more effective, efficient, and secure, we can develop confidence in the sustainability of their implementation.

Table 2:   Expected Benefits of this Research

Beneficiary	Benefit	Supporting Research Outcome
Internal Auditors and Information Security Managers	*a fact-based case to recommend the implementation and rigorous treatment of change and access controls*	*by providing them with a case demonstrating the foundational nature of these controls*
IT Managers and Administrators	*a better understanding of the pitfalls associated with decisions to bypass these controls*	*by making them aware of the long-term unintended and unanticipated negative impacts of their decisions on performance*
IT and Business Executives	*an enhanced confidence in showing a return on investment on the implementation of change and access controls*	*by illuminating the relationship between these controls and improved performance, which leads to a higher business value*

2 Methodological Background

Our research uses a technique called system dynamics—a method for modeling and analyzing the holistic nature of complex problems as they evolve over time [Sterman 2000]. System dynamics has been used to gain insight into some of the most challenging strategy questions facing businesses and government for several decades. The Franz Edelman Prize for excellence in management was given in 2001 to a team at General Motors who used system dynamics to develop a successful strategy for launch of the OnStar System [Barabba 2002]. System dynamics is particularly useful for gaining insight into difficult management situations in which best efforts to solve a problem actually make it worse. Real problematic situations in which system dynamics helps to create clarity include the following [Sterman 2000]:

- Efforts to build new roads to alleviate traffic congestion only result in increased congestion.

- Use of cheaper drugs pushes costs up, not down.

- Lowering the nicotine in cigarettes, supposedly to the benefit of smoker's health, only results in people smoking more cigarettes and taking longer, deeper drags to meet their nicotine needs.

- Levee and dam construction to control floods leads to more severe flooding by preventing the natural dissipation of excess water in flood plains.

- Applying more resources to incident response to handle a high workload takes resources from proactive management activities and increases the incident workload.

Here system dynamics targets problematic behavior associated with business operations in general and IT management in particular. Intuitive solutions to problems in this area often reduce the problem in the short term, but make it much worse in the long term. System dynamics is a valuable analysis tool for gaining insight into solutions that are effective over the long term and demonstrating their benefits.

A powerful tenet of system dynamics is that the dynamic complexity of problematic behavior is captured by the underlying feedback structure of that behavior. So we decompose the causal structure of the problematic behavior into its feedback loops to understand which loop is strongest (i.e., which loop's influence on behavior dominates all others) at particular points through time. We can then thoroughly understand and communicate the nature of the problematic behavior and the benefits of alternative mitigations.

System dynamics model boundaries are drawn so that all the enterprise elements necessary to generate and understand problematic behavior are contained within them. This approach encourages the inclusion of soft factors in the model, such as

policy, procedural, administrative, or cultural factors along with hard, strictly technical factors. The exclusion of soft factors essentially treats their influence as negligible when in fact it is frequently significant. This endogenous viewpoint helps show the benefits of mitigations to the problematic behavior that are often overlooked by low performers, partly due to their narrow focus on technical solutions to resolve problems.

We rely on system dynamics as a tool to help test the effect of strategies for improving the performance of IT management. In some sense the simulation of the model will help predict the effect of these strategies. But what is the nature of the types of predictions that system dynamics facilitates? Dennis Meadows offers a concise answer by categorizing outputs from models [Meadows 1974]:

1. absolute and precise predictions (Exactly when and where will the next cyber attack take place?)

2. conditional precise predictions (If a cyber attack occurs, how much will it cost my organization?)

3. conditional imprecise projections of dynamic behavior modes (If I adopt IT change management controls, will my business's performance be better than it would have been otherwise?)

4. summary and communication of current trends, relationships, or constraints that may influence the future behavior of the system (If the current trends in distributed denial-of-service attacks continue, what effect will this have on my business over then next five years?)

5. philosophical explorations of the logical consequences of a set of assumptions, without any necessary regard for the real-world accuracy or usefulness of those assumptions (How would genetic experimentation that allows development of human telepathic abilities affect a business's exposure to insider threat?)

The model we develop, and system dynamics models in general, provide information of the third sort. Meadows explains further that "this level of knowledge is less satisfactory than a perfect, precise prediction would be, but it is still a significant advance over the level of understanding permitted by current mental models."

2.1 NOTATION

In graphic representations of the model we describe, signed arrows represent the system interactions, where the sign indicates the pair-wise influence of the variable at the source of the arrow on the variable at the target of the arrow:

- Roughly, an arrow labeled with a + indicates that the value of the source and target variables move in the same direction.[4]

- Roughly, an arrow labeled with a - indicates that the value of the source and target variables move in the opposite direction.[5]

We can illustrate the above definitions using the influence diagram shown in Figure 1, which represents a very simple room heating system. A positive influence is indicated by the arrow from *rate of heat input* to *room temperature*. At a particular thermostat setting, as the rate of heat input increases (or decreases), then the temperature of the room increases (or decreases) above (or below) what it would have been. A negative influence is indicated by the arrow in the other direction. As the room temperature increases (or decreases), the rate of heat input decreases (or increases) below (or above) what it would have been, as would be expected by a room heating system.

thermostat
setting

rate of heat
input

+

room
temperature

-
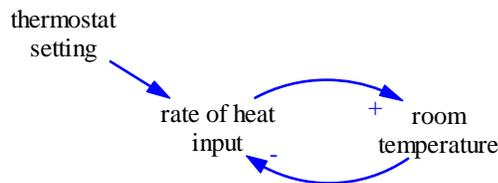
Figure 1:   A Simple Feedback Loop

As mentioned previously, dynamically complex problems can often be best understood in terms of the feedback loops underlying those problems. There are two types of feedback loops: balancing and reinforcing. Balancing loops describe aspects of the system that oppose change, seeking to drive organizational variables to some goal state. Reinforcing loops describe system aspects that tend to drive variable values consistently upward or consistently downward. The polarity of a feedback loop is determined by "multiplying" the signs along the path of the loop. Balancing loops have negative polarity and reinforcing loops have positive polarity.

Figure 1 depicts a balancing loop that seeks to move the room temperature to the thermostat setting. This system is balancing as shown by the odd number of negative signs along its path. The goal state is a room temperature equal to the thermostat setting. In general, balancing loops describe aspects that oppose change, and usually involve self-regulation through adaptation to external influences.

---

[4]   More formally, a positive (+) influence indicates that if the value of the source variable increases, then the value of the target variable increases above what it would otherwise have been, all other things being equal. And if the value of the source variable decreases, then the value of the target variable decreases below what it would otherwise have been, all other things being equal.

[5]   More formally, a negative (-) influence indicates that if the value of the source variable increases, then the value of the target variable decreases below what it would otherwise have been, all other things being equal. And, if the value of the source variable decreases, then the value of the target variable increases above what it would otherwise have been, all other things being equal.

Figure 2 shows a more interesting example in the domain of project management. Figure 2a depicts one approach an organization may adopt to try to put a project that is behind schedule back on track: having its employees work overtime. The closed form in Figure 2b shows the corresponding balancing feedback loop that characterizes the goal of the approach as moving the project to the state of being on schedule.
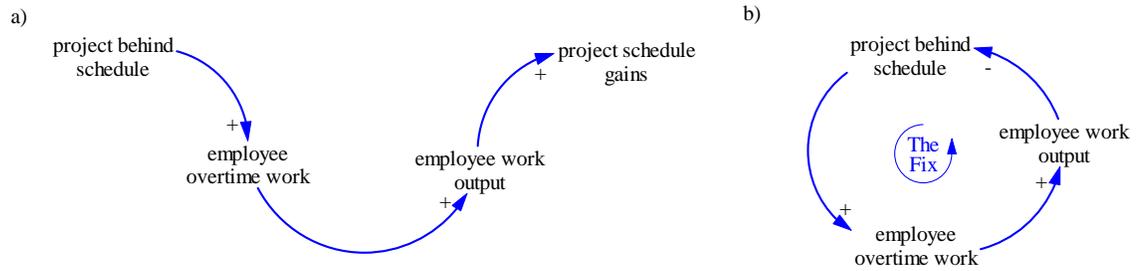
Figure 2: (a) Project Management—Desire to Use Overtime to Correct Schedule; (b) Closed-Loop Representation Showing (Balancing) Feedback to Improve Progress

Figure 3 shows that the project-management behavior described above is subject to a reinforcing feedback loop in which overtime in the long term leads to employee burnout, lower quality of work, and the need to rework defective artifacts. The longer this goes on the further the project gets behind schedule because of the increasing amount of rework. This combines with the previous balancing feedback loop, where the balancing loop dominates in the near term with the reinforcing loop taking over with increasing amounts of employee overtime and burnout. This type of thinking about the feedback structure of systems and about which feedback loops dominate at different periods in time is characteristic of system dynamics modeling and analysis.

The reinforcing loop is shown mostly in red but it shares part of the influence path of the blue balancing loop from *project behind schedule* to *employee overtime work*. The reinforcing nature of the feedback loop is evident from the even number of negative signs along its path.[6] Reinforcing loops may help explain explosive growth or implosive collapse of a system.

---

[6] Feedback loops that have no negative signs along the influence path have positive polarity and thus are reinforcing loops.
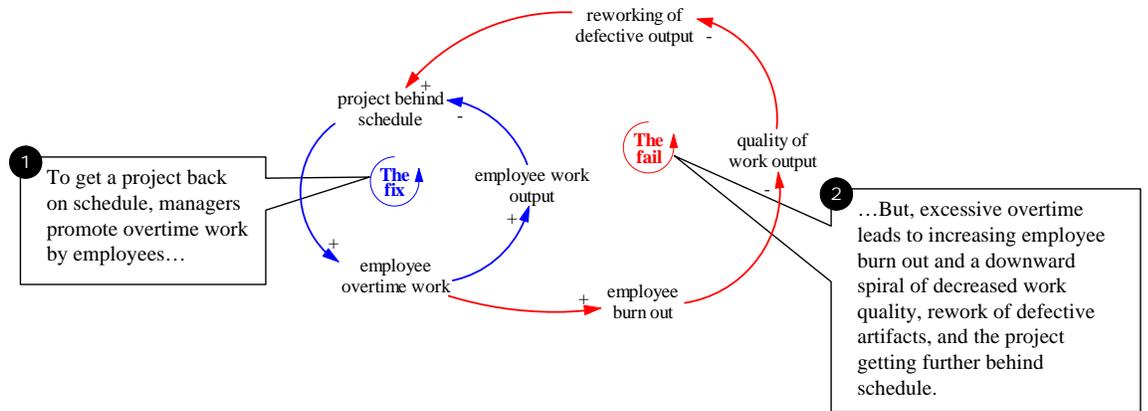
Figure 3:   Unintended Burnout due to Overtime

3 System Dynamics Model

We hypothesize that the primary reasons for low IT management performance are

- an overly reactionary approach to operational problems
- a tendency to erode change and access controls over time to expedite the fire-fighting needed to maintain business operations

The result is a patchwork of unofficial and undocumented workarounds, supporting a patchwork of increasingly unstable and undocumented software and systems that continue to degrade with time. Moreover, the combination of fragmenting processes and IT systems serves to undermine the organization's ability to understand and control the operational environment, leading to a downward spiral of ever-increasing operational problems.

The appendices to this paper provide a summary of primary assumptions that the model makes and a comprehensive graphical overview of the system dynamics model that we have developed, which is described more fully by Moore.[7] This section presents the essential elements of that model. We first characterize the nature of change and access controls. We then present the basic stock and flow structure of the model to characterize the primary underlying accumulations and flows that are relevant to the low performer's problematic behaviors. Using this structure as a basis, we then present a high-level view of a low performer's decision making in terms of the primary feedback loops. For traceability, the feedback loops presented here are labeled identically to those in the full stock and flow model described in Appendix B.

Simulation of the model allows comparison of results with known historical behavior of the low performer. Once we have confidence that the simulation model accurately captures the low-performer problematic behavior, we will be in a position to determine the benefit of strategies for improved business performance and security, including more rigorous change and access controls.

3.1 NATURE OF CHANGE AND ACCESS CONTROLS

Change and access management processes are often viewed as a series of tasks to be accomplished. This, however, is only a partial description of what a process truly is. Garvin explains that a process is made up of two components: physical and behavioral [Garvin 1995]. The physical component—which is tangible and therefore gets most of the attention—is defined as a work process that consists of a se-

---

[7] Moore, A.P. & Antao, R. *System Dynamics Modeling and Analysis of IT Management Controls in Context.* Pittsburgh, PA: Software Engineering Institute Special Report, 2005. Not publicly available.

quence of linked, interdependent activities, which, taken together, transform inputs to outputs [Garvin 1995].

Take the change management (CM) process, for instance. We can view the physical component as a work process that takes requests for change (RFC) as inputs and produces successfully implemented changes that are documented. Between the input and output phases the requested change progresses through a number of interdependent tasks such as change planning, authorization, testing, documentation, and implementation, as shown in Figure 4 [Behr 2004].
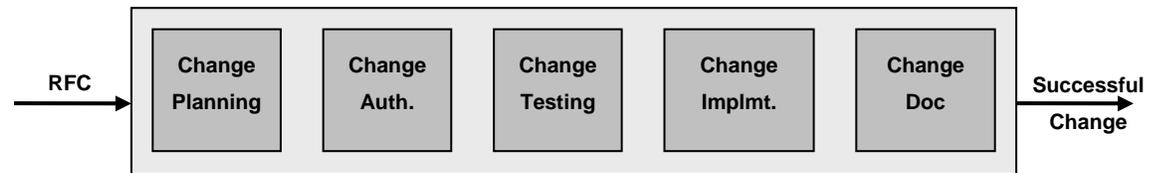
Figure 4:   The Physical Components of the Change Management Process

The behavioral component, on the other hand, is an underlying pattern of decision making, communication, and learning that is deeply embedded and recurrent within an organization. Behavioral components have no independent existence apart from the work processes in which they appear. Nevertheless, these components profoundly affect the form, substance, and character of activities by shaping how they are carried out. To truly understand the functioning of an organization's IT operations process with respect to change and access management, we must consider both the physical and behavioral aspects of this process.

3.2 BASIC STOCK AND FLOW INFRASTRUCTURE

A quantitative system dynamics model refines and describes the relationships in the qualitative system dynamics model using mathematical equations. This is done by adding two new concepts to the modeling notation: stocks and flows.

1. Stocks represent accumulations of physical or non-physical quantities and flows represent the movement of these quantities between stocks. Stocks are depicted as named boxes within the model.

2. Flows are depicted as double-lined arrows between the stocks with a named valve symbol indicating the name of the flow. Flows that come from (or go to) a cloud symbol indicate that the stock from which the flow originates (or to which the flow goes) is outside the scope of the model.

The next subsection describes the stock and flow infrastructure of our system dynamics model. The rest of the section then describes the feedback loops that characterize IT management decision making and operations in terms of the stock and flow infrastructure.

3.2.1 The Service View

Figure 5 shows the service view of the stock and flow model. The *Critical operational services* stock includes those services that are currently running and fully operational. Services can be upgraded in a planned way or they can fail and be fixed in an unplanned way. The *Planned upgrades* stock includes those services that have been taken offline for some period of time to install the upgrade.[8] Upgrades are the result of planned changes of service *artifacts*, which will be described in the next section.

Service failures exhibit themselves as degraded operations or non-operation. They may be caused by

- malicious individuals wishing to do the organization harm, either internal or external to that organization

- stresses imposed on the services due to authorized use by legitimate users

- failures due to the aging of (hardware) artifacts that support those services

---

[8] We include system maintenance in the class of service upgrades if it requires the service to be brought down for some period of time.
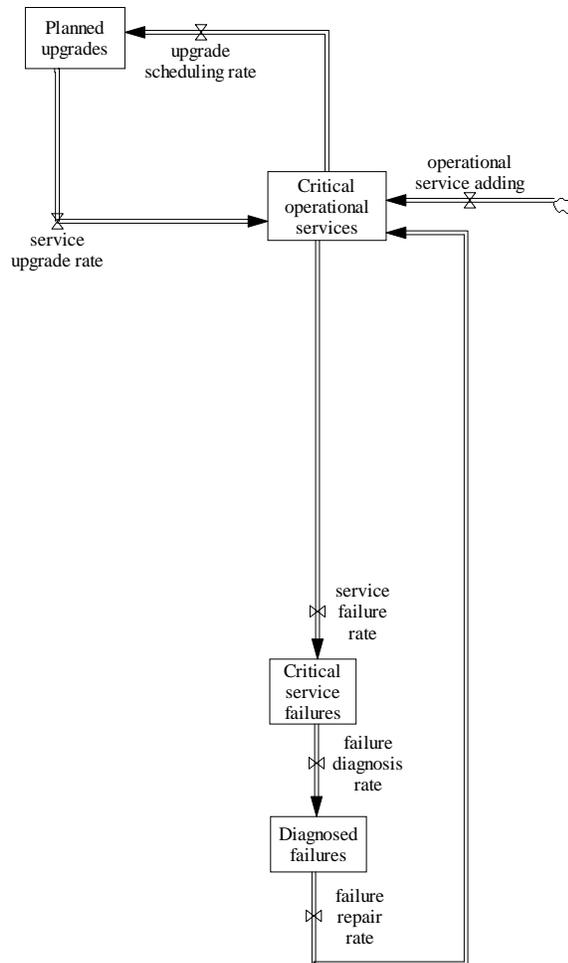
Figure 5:   Service Flows

The stock and flow model separates failure diagnosis and failure repair since the rate of these two activities have different influencing factors. The *Critical service failures* stock contains those services that have failed with the reason for the failure not yet determined. The *Diagnosed failures* stock contains those failed services that have been diagnosed. Of course, in real operations, a failed service is likely to be brought back up in degraded mode while the cause of the failure is diagnosed and repaired. In this model, such a failed service would not be in the stock of *Critical operational services* until that repair has been made.

3.2.2   The Artifact View

Figure 6 shows the basic flows of the artifact view of the model. The artifact view is the static (developmental) counterpart of the dynamic (operational) service view. Flows in the artifact and the service views march in synchronized step. Upgrade scheduling in the service view leads to *Planned work to do* in the artifact view. Planned work may involve creating new artifacts, changing existing artifacts, or retiring old artifacts.

Operational services include artifacts that may either be classified (grossly) as *Reliable artifacts* or *Unreliable artifacts*. The reliability of artifacts produced as a result of planned work depends on the *planned change success rate*. The planned change failure rate is simply one minus the planned change success rate. Analogous to the failure of services, reliable artifacts may become unreliable, via the *losing artifact reliability* flow, for the following reasons:

- vulnerabilities discovered that may be exploited by malicious individuals

- new unforeseen uses of the artifacts beyond that for which they were designed

- aging of (hardware) artifacts

Unreliable artifacts eventually lead to *Problem work to do*, which is identified when a service fails and the reason for that failure is determined. The failure diagnosis identifies the (previously unknown) unreliable artifacts as the culprit in the failure. Subsequent repair of the problem leads to bringing the service back into operation. Of course, repair work may not itself be perfect so some of the repaired artifacts may remain unreliable, indicating the potential for additional future service failures.

A major aspect of our hypothesis about the cause of low performance in IT management is that an overly reactive approach that focuses on emergency repair work to keep IT services up and running results in a fragile IT environment that is subject to high change failure rates. Fragile IT environments are built on fragile artifacts. Fragile artifacts are those artifacts that may operate reasonably well in operation, but when changes are made to other artifacts that depend on them, the chance of failure is high. As described in *The Visible Ops Handbook*, fragile artifacts generate much firefighting and need to be identified and handled with care [Behr 2004].
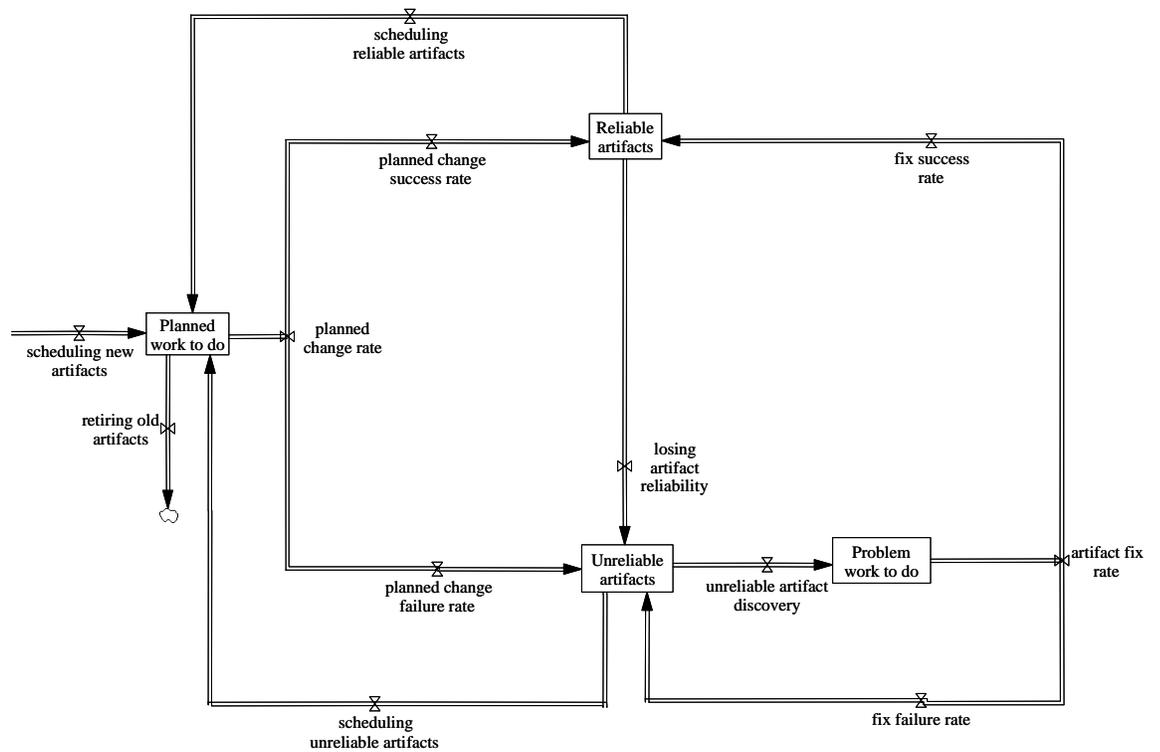
Figure 6:   Basic Artifact Flows

A high-leverage fundamental solution for IT management suffering low perform-
ance, then, should be to find and fix those fragile artifacts that are embedded in
their IT infrastructure. Figure 7 depicts an extension to the stock and flow infra-
structure of our system dynamics model. Three stocks of artifacts are added: *Non-
fragile artifacts*, *Undiscovered fragile artifacts*, and *Discovered fragile artifacts*.[9]
Nonfragile artifacts become fragile as a result of changes to the system, particularly
problem fixes. A fragile artifact is typically discovered as a result of the diagnosis
of service failures, particularly if the artifact is the regular cause of the failure.
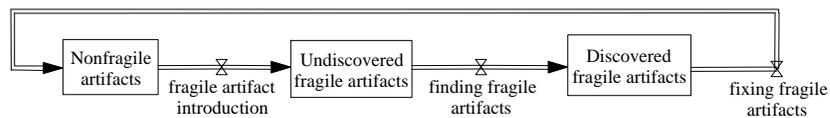
Figure 7:   Flows Involving Artifact Fragility

---

9   Fragile artifacts are different from Unreliable artifacts since fragile artifacts may operate reasonably
     well in an unchanging environment. It is only when a fragile artifact's environment is changed that the
     fragile artifact may cause a problem. Unreliable artifacts cause problems due to the stress of opera-
     tions, whereas fragile artifacts cause problems due to the stress of change. Of course, an artifact may
     be both unreliable and fragile.

3.2.3 The Personnel View

Figure 8 depicts the personnel view of the model. There are only two types of personnel considered in the model: Planned-change personnel and problem-repair personnel. Planned-change personnel are responsible for planned changes to artifacts that happen as a result of planned upgrade to services. Problem-repair personnel, on the other hand, diagnose and fix unreliable artifacts discovered as a result of service failures.
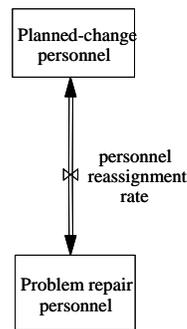
Figure 8: Personnel Flows

Personnel may be reassigned in either direction—planned-change personnel may be reassigned to problem (service failure) work or problem-repair personnel may be reassigned to planned work (service upgrades). However, it is not within the scope of the model to include facilities for hiring additional personnel. While this is certainly an important option in real-world management, all organizations operate under constraints that do not always permit hiring additional personnel even if that would help alleviate their problem. The point of the current model is to see how well organizations can do with the staff that they have on hand.

3.3 FEEDBACK STRUCTURES

We now present the primary feedback loops of the stock and flow model presented in Appendix B. As mentioned, we label the feedback loops identically to those in the full stock and flow model for traceability. We also use boxes to highlight those stocks that are part of the stock and flow infrastructure presented in the last section. Colors used in this and subsequent causal loop diagrams are used to help distinguish the different feedback loops.

Two archetypes are particularly relevant for the models that we develop in this paper: the Fixes that Fail archetype and the Shifting the Burden archetype. We use these archetypes as the basis for describing the model.[10]

---

[10] These archetypes are special cases of the Out of Control archetype as described by Wolstenholme [Wolstenholme 2003].

3.3.1 IT Management "Fixes that Fail"

Senge describes the generic Fixes that Fail archetype very simply as follows:

A fix, effective in the short term, has unforeseen long-term consequences which may require even more use of the same fix [Senge 1990].

This archetype, which is shown in Figure 9, contains one balancing loop—the "fix"—that decreases the problem in the short term. The unintended consequences of the fix often take longer to occur and increase the problem in a self-reinforcing way in the long term. The project-management influence diagram that we characterized previously in Figure 3 is an example of a Fixes that Fail archetype, where the fix is the overtime work to get back on schedule—the balancing loop—and the unintended consequence is the burnout due to excessive overtime—the reinforcing loop.

1 An organization attempts to fix a problem, but the fix is effective only in the short term, because…

2 …Unintended and unexpected consequences of the fix increase the problem leading to application of even more of the same fix. This results in a worsening problem situation with no effective long-term remedy.
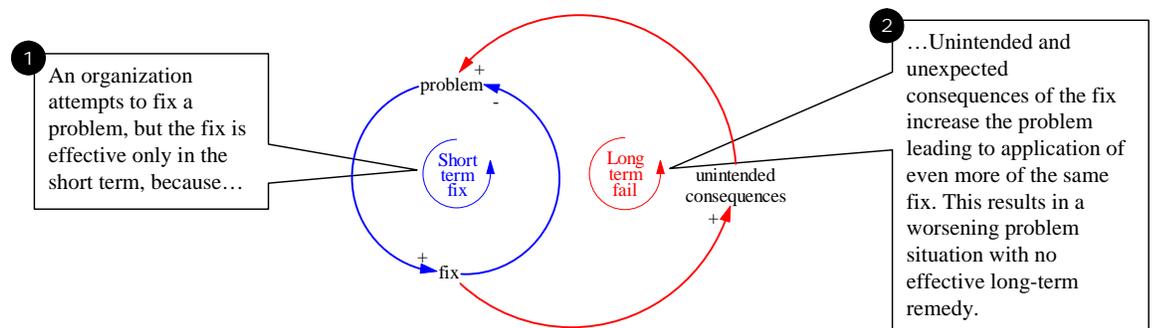
Figure 9:   Fixes that Fail Archetype

We hypothesize that there are four main approaches that low performers use to manage IT operations. We hypothesize that these actions bring about the majority of problems for IT management low performance:

1.   relaxing IT change testing quality

2.   relaxing IT change documentation quality

3.   relaxing access controls on IT operations and development staff

4.   shifting personnel to problem work

These actions may occur more by reflex in the heat of the moment rather than as an explicit action by management. Nevertheless, they are all intended to improve system availability and lessen the work pressure on IT operations staff.

IT change includes either planned change or unplanned change. We refer to unplanned changes as problem fixes. Figure 10 illustrates the first three of the above approaches and the unintended consequences that they bring:

- Loop *B1* reduces the problem fix testing level with the unintended fix quality degradation of loop *R1*: Fix testing can encompass a large percentage of the effort and time associated with repairing failed services and bringing them

back online. However, decreased fix testing degrades the quality of fixes to service problems which, in turn, degrades the reliability of system artifacts.

- Loop *B2* reduces fix documentation with the unintended fix diagnosis degradation of loop *R2*: Fix documentation may also take a fair amount of time to do properly. However, degraded change documentation leads to difficulty diagnosing IT problems that involve the poorly documented changes. Diagnosis difficulties result in longer repair times.

- Loop *B3* reduces the controls associated with IT Ops staff access to artifacts with the unintended artifact corruption of loop *R3*: Relaxed access controls give problem-repair personnel easy access to the system with no time wasted waiting for the right kind of authorization. This allows personnel to understand the root cause of failures and get full operations back up and running as quickly as possible. However, as access controls are relaxed, the operations staff gradually loses control over exactly who has made what changes to the system. Even worse, people start making changes completely unrelated to system failures. These effects result in corruption of system artifacts and degrading of their reliability.

In summary, the organizational responses described by loops B1 through B3 intend to get failed services back up and running as soon as possible, but the over-reliance on these methods can, in the longer term, result in a downward spiral toward more and more downtime as seen by the reinforcing loops R1 through R3.
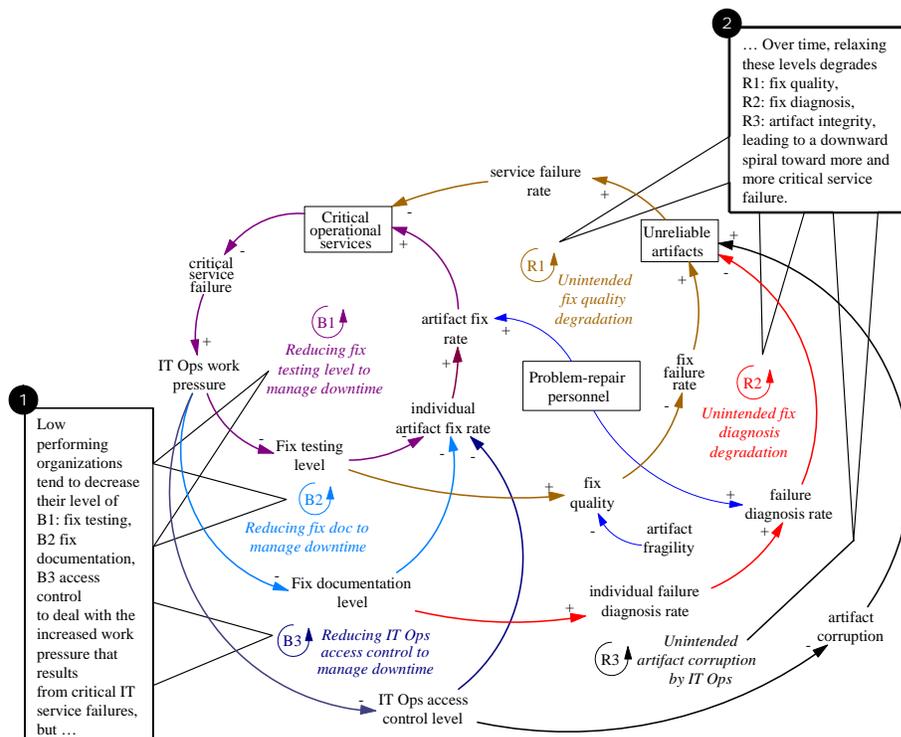
Figure 10: Relaxing Change and Access Controls to Manage Downtime

Figure 11 illustrates the fourth response of low-performing organizations to IT Ops work pressure: shifting personnel to problem-repair work. This response, depicted by loop *B4*, is a natural and often useful reaction for increasing failure repair rate and bringing services back up and running as soon as possible. As shown in the figure,

this response leads to reductions in planned-change personnel and a number of un-intended consequences, which parallel the unintended consequences seen in Figure 10:[11]

- Unintended planned change quality degradation (loop *R4*): Shortages in planned-change personnel can result in relaxed planned change testing due to increased work pressure on the development staff. As in the case of relaxed fix testing by IT operations, this leads to degraded artifact quality.

- Unintended planned change documentation quality degradation (loop *R5*): Development staff work pressure can also result in lower levels of planned change documentation. This result can inhibit service failure diagnosis and the repair of unreliable artifacts.

- Unintended artifact corruption by IT development staff (loop *R6*): Finally, development staff work pressure can result in relaxing access controls on IT development staff. As in the case for the IT operations staff, this response gives planned-change personnel easy access to the system with no time wasted waiting for the right kind of authorization. However, as access controls are re-laxed, the organization gradually loses control over exactly who has made what changes to the system. Moreover, changes made by the development staff start to clash with changes made by the operations staff to fix service problems. These effects result in corruption of system artifacts and degrading of their reliability.

In summary, shifting personnel from planned work to problem-repair work can, in the longer term, add to the downward spiral of the organization toward more and more downtime, requiring even more personnel shifting to problem-repair work.

---

[11] There are also balancing loops in the IT development domain that parallel the "fixes" associated with loops *B1* through *B3* in Figure 10. For simplicity, these balancing loops are not shown.
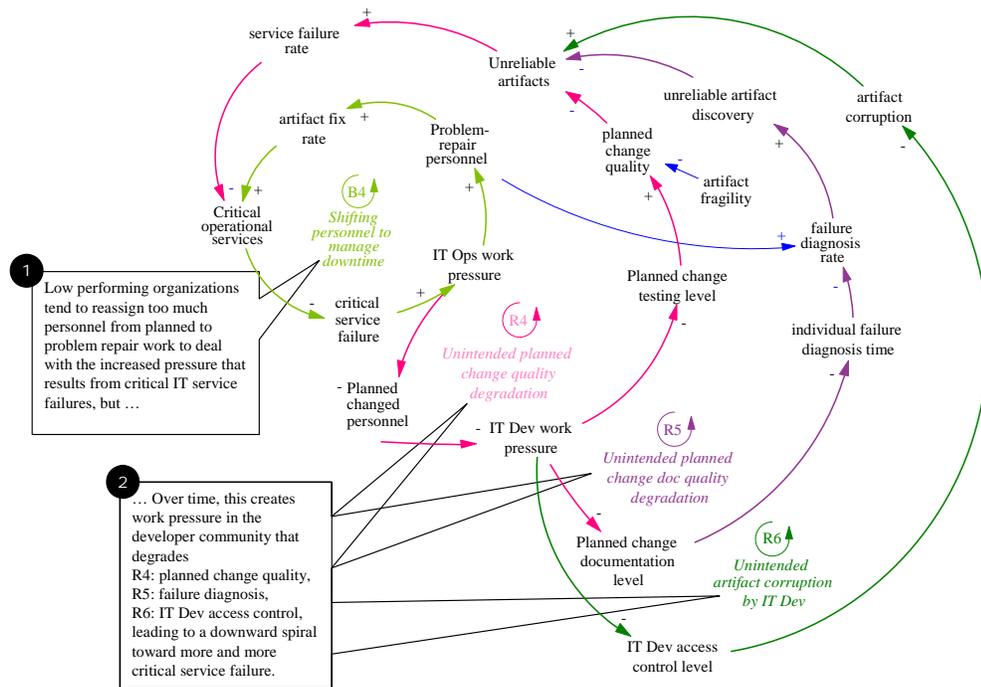
Figure 11: Shifting Planned-Change Personnel to Problem Management

3.3.2 IT Management "Shifting the Burden"

Senge defines the Shifting the Burden archetype as follows:

An underlying problem generates symptoms that demand attention. But the underlying problem is difficult for people to address, either because it is obscure or costly to confront. So people "shift the burden" of their problem to other solutions—well-intentioned, easy fixes which seem extremely efficient. Unfortunately, the easier "solutions" only ameliorate the symptoms; they leave the underlying problem unaltered. The underlying problem grows worse, unnoticed because the symptoms apparently clear up, and the system loses whatever abilities it had to solve the underlying problem [Senge 1990].

Figure 12 depicts the Shifting the Burden archetype. The balancing feedback loop at the bottom of the figure represents the attempt to address a problem symptom by an organization as an easy fix to put the organization back on track, instead of addressing underlying root causes using a fundamental solution (top loop). Symptomatic solutions often result in a reinforcing loop, shown on the left side of the figure, in which the symptomatic solution can cause the capability for fundamental solutions to atrophy gradually over time. For example, in the project-management problem, described in the previous section

- The symptomatic solution was to engage workers in overtime to put their project back on schedule.

- The fundamental solution might have been to increase the hiring rate.

- The fundamental solution was gradually degraded by over-application of the symptomatic solution because burned-out workers often quit, leading to damaged organizational reputation and difficulty hiring.
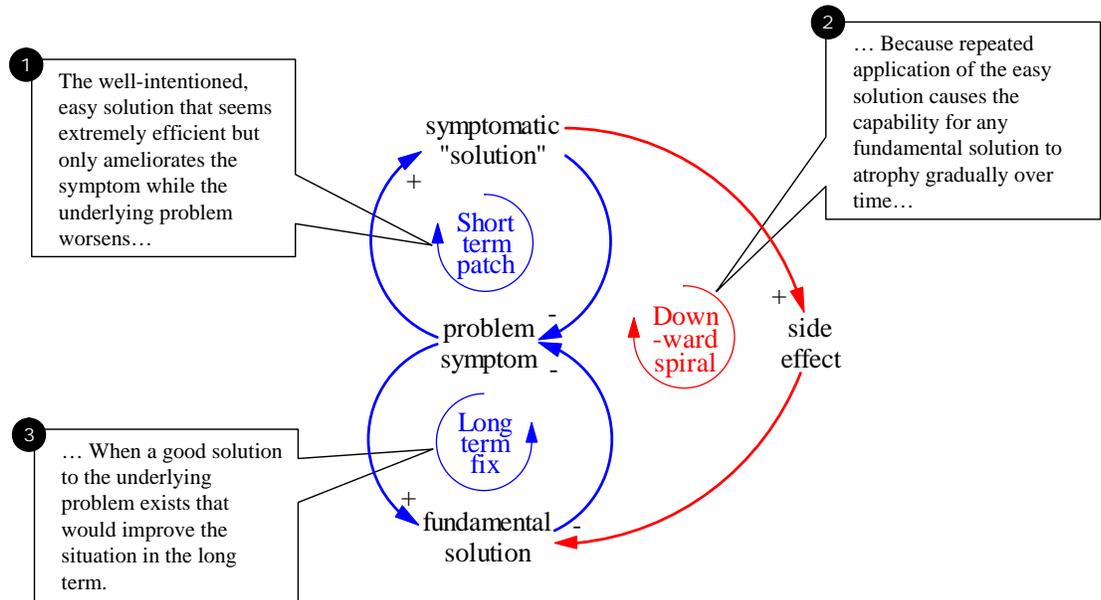
Figure 12:   Shifting the Burden Archetype

Figure 13 shows two classes of solutions available to IT managers to handle the problem of critical service failure: the symptomatic and fundamental solutions. The IT manager must decide how to split organizational resources between reactive and proactive activities. Symptomatic solutions are typically reactive in nature. The feedback loop labeled B4 in Figure 13 is an example of a symptomatic solution to the problem of service failure. This is the same loop labeled B4 depicted in Figure 11. Shifting personnel to problem work is a natural managerial action to excessive downtime which can be effective in the short term. However, low performers often move too many of their resources to incident response at the first sign of problems.

Fundamental solutions to excessive downtime identify strategies for the evolution of information systems toward higher system availability in the long term. With increased identification of high-confidence solutions to availability problems comes increased implementation of these proactive solutions leading to higher availability over the long term. Such fundamental solutions have been very successful in practice [Stern 2001].

The feedback loop labeled *B5* in Figure 13 poses a particular fundamental solution to the problem of excessive downtime. It involves finding and fixing fragile artifacts. A system is fragile if, when subjected to a change, the system is highly likely

to fail. We refer to this as a *change failure*. A fragile system is one that is highly dependent on fragile artifacts. Thus, finding and fixing fragile artifacts reduces system fragility and thus increases the change success rate given the same amount of change testing.

While fundamental solutions are important to the long-term health of organizational operations, clearly some immediate relief must go to the problem of service failure. However, as shown in the *R7a* loop of Figure 13, too much focus on reactive activities that reassign personnel from planned work to problem work takes resources away from finding and fixing fragile artifacts. Loop *R7b* shows that continual patching of IT problems increases artifact and system fragility and leads, over time, to decreased control and understanding of the IT operational environment. The result is lowered change success rate due to higher system fragility and even more system failure. This worsening situation is characteristic of the Shifting the Burden archetype and the downward spiral of the low performer.

Figure 13: Reactivity Degrading Long-Term Availability

4   Simulation Results

This section describes preliminary simulation results obtained by executing the model described in the last section. The behavior of the model is based on a set of functions that have the general form "Effect of X on Y." The inputs and outputs of these functions are normalized so that

- the input value is the dimensionless ratio of the X to a normal value for X and
- the output is a dimensionless effect modifying the normal value for Y

That is, for function f which describes the effect of X on Y, Y=normal Y*f(X/normal X) as described by Sterman [Sterman 2000]. Normal values across the model are specified with respect to a *user standard service failure*, intended to be the maximum level of failure users will find generally acceptable.[12]

Our results are described with respect to a model equilibrium in which the inflows of all stocks equal their outflows. Such equilibrium ensures that all stocks remain at a constant level. In equilibrium, it is relatively easy to validate and to experiment with a model since the analyst can more readily determine how small changes in input affect the overall behavior through simulation. Any change in behavior (as seen in the time graphs) can be attributed to that change and only that change. It is analogous in scientific experiments to keeping all variables constant except the ones being studied.

The rest of this section describes how the model responds to a perturbation of its input: the step increase in *vulnerabilities discovered* in organizational systems. These vulnerabilities could arise from exploits discovered within operating artifacts or from artifact aging. Intuitively, this increase might be attributed to an expanding hacker community that is dedicated to finding and exploiting vulnerabilities in current technologies.

4.1   MODEL RESPONSE TO INPUT PERTURBATION

The following organizational responses to the new input are tested:
- The organization executes business as usual with little to no commitment behind change controls (respectively, access controls and staffing of planned work). As work pressures rise, the organization reduces its change controls (respectively, access controls, and staffing of planned work) to more quickly

---

[12]   The *user standard service failure* parameter is analogous to a customer-driven requirement for reliable system operation.

implement emergency fixes. Reduction in change controls constitutes a reduction in change testing and/or change documentation.

- The organization closely adheres to its change controls (respectively, access controls, and staffing of planned work) with the hopes that higher quality fixes and continuance of planned work will pay returns in the long run.

Figure 14 shows the *critical service failure* that results over time with a 50% rise in vulnerabilities at the fourth week in the simulation. The baseline run, displayed in blue and labeled 1, shows the system to be in equilibrium with respect to the level of failure. The rest of the runs, labeled 2 through 9, show the critical service failure with various combinations of policies:

- Change control

  - C: committed to change control policy
  - nC: relaxes enforcement of change control policy when need arises

- Access control

  - A: committed to access control policy
  - nA: relaxes enforcement of access control policy when need arises

- Shifting personnel from planned to emergency repair work:

  - F: flexible policy regarding moving people to unplanned work
  - nF: ensures minimum level of staffing of planned work

The eight combinations of the above policies are reflected in the eight runs (in addition to the baseline) displayed in the figure.

We make the following observations about the above runs:

- The use of change and access controls is subject to a worse-before-better performance. There are some early throughput gains from not using these controls but the long-term advantages of using them outweigh their short-term disadvantages. Managers must be aware of the short-term disadvantages so they can last through them to accrue the long-term advantages.

- Shifting personnel from planned work to problem-repair work to manage downtime can work in the short term, but at long-term costs that can overwhelm an organization's ability to cope. Some discriminate shifting of personnel may be needed to achieve short-term goals, but care must be taken not to sacrifice long-term performance. Future work will test the tradeoffs inherent in this approach.

The better performance through the use of IT controls assumes that an organization has limited resources to put into problem management. This is after all, a practical business reality.
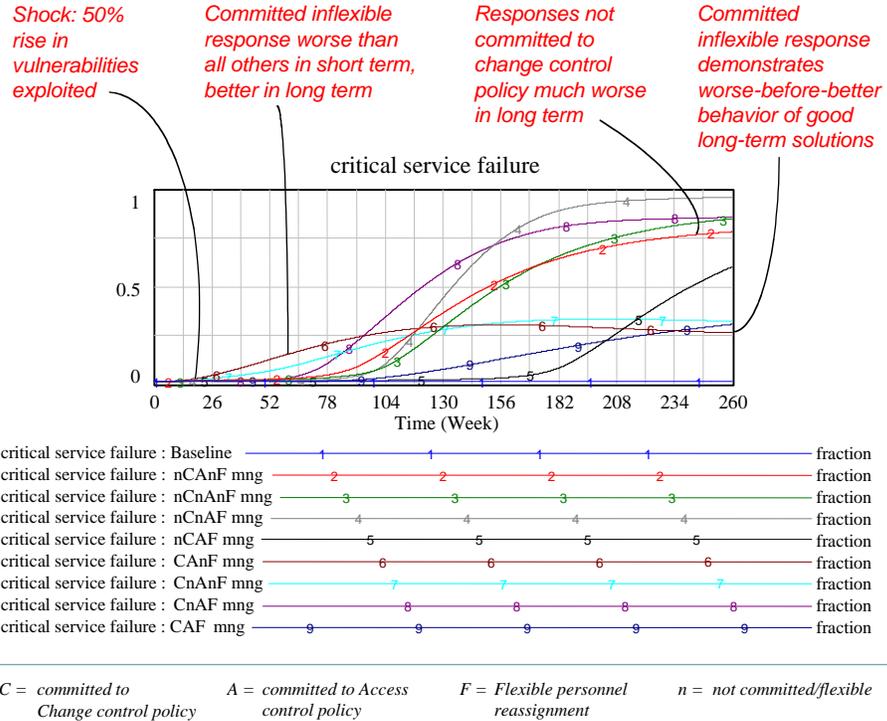
Figure 14: *Results from Increasing Vulnerability Discovery by 50% for Critical Service Failure*

Figure 15 shows the results of increasing vulnerability introduction in the model by 50% with respect to two performance measures: percentage of unplanned work and percentage of change successes.[13] It is not too surprising that Figure 15a shows that percentage of unplanned work grows faster and higher in the case (F) where personnel can be shifted from planned work to unplanned work (i.e., runs 4, 5, 8, and 9). In these cases, it takes from a year to 18 months for almost all of the planned work personnel to be transferred.

---

[13] Percentage of unplanned work is defined within the model as the ratio of *artifact fix rate* to the total change rate. The total change rate is the sum of the *planned change rate* and the *artifact fix rate*. Percentage of change successes is defined in the model as the ratio of the sum of the *fix success rate* and the *planned change success* rate to the total change rate.

percentage of unplanned work

percentage of unplanned work : Baseline	1	1	1	fraction
percentage of unplanned work : nCAnF mng	2	2	2	fraction
percentage of unplanned work : nCnAnF mng	3	3	3	fraction
percentage of unplanned work : nCnAF mng	4	4	4	4	fraction
percentage of unplanned work : nCAF mng	5	5	5	5	fraction
percentage of unplanned work : CAnF mng	6	6	6	fraction
percentage of unplanned work : CnAnF mng	7	7	7	fraction
percentage of unplanned work : CnAF mng	8	8	8	fraction
percentage of unplanned work : CAF mng	9	9	9	fraction

b)

percentage of change successes

percentage of change successes : Baseline	1	1	1	fraction
percentage of change successes : nCAnF mng	2	2	2	fraction
percentage of change successes : nCnAnF mng	3	3	3	fraction
percentage of change successes : nCnAF mng	4	4	4	fraction
percentage of change successes : nCAF mng	5	5	5	5	fraction
percentage of change successes : CAnF mng	6	6	6	fraction
percentage of change successes : CnAnF mng	7	7	7	fraction
percentage of change successes : CnAF mng	8	8	8	fraction
percentage of change successes : CAF mng	9	9	9	fraction

Key: *C = committed to Change control policy*   *A = committed to Access control policy*   *F = Flexible personnel reassignment*   *n = not committed/flexible*

Figure 15: Results from Increasing Vulnerability Discovery by 50% for a) percentage of unplanned work and b) percentage of change successes

The remaining runs of Figure 15a are somewhat more interesting. Operations that enforce change control policy (i.e., runs 6 and 7) have much better percentage of unplanned work than operations that do not (i.e., runs 2 and 3). This is primarily due to the fact that non-commitment to change controls leads to high levels of service failure that inhibits planned change work.[14] Similarly, the enforcement of

---

[14] We assume that failed services cannot be scheduled for planned work—they must be returned to the operational state before planned changes can commence.

access controls leads to higher planned to an unplanned work ratio. In general, planned work can proceed in a more straightforward and scheduled way when operational services are not regularly interrupted with failures.

Figure 15b shows that in terms of change success operations committed to change controls (i.e., runs 6 through 9) outperform operations that are not so committed (i.e., runs 2 through 5). Again, this is not too surprising since operations committed to change controls maintain the quality of both planned change testing and problem fix testing necessary to promote change success. The remaining runs show that operations committed to full staffing of planned work (i.e., runs 6 and 7) perform better than operations not so committed (i.e., runs 8 and 9). This is primarily due to the increased fragility that results from pulling people from planned work to increase levels of patching to get services up and running. Over time, the operational environment erodes with such an emphasis on patching, making it increasingly difficult to implement successful changes.

4.2 TESTING DIFFERENT LEVELS OF CHANGE CONTROL

The analysis performed in the previous section assumes a normal change control level of 0.5 in the range zero to one. Lack of commitment to change control can result in reduced change control (i.e., relaxed change/fix testing or documentation) but we did not test operational behavior for increased change control. Figure 16 verifies that lower levels of change control do lead to greater critical service failure in the model.
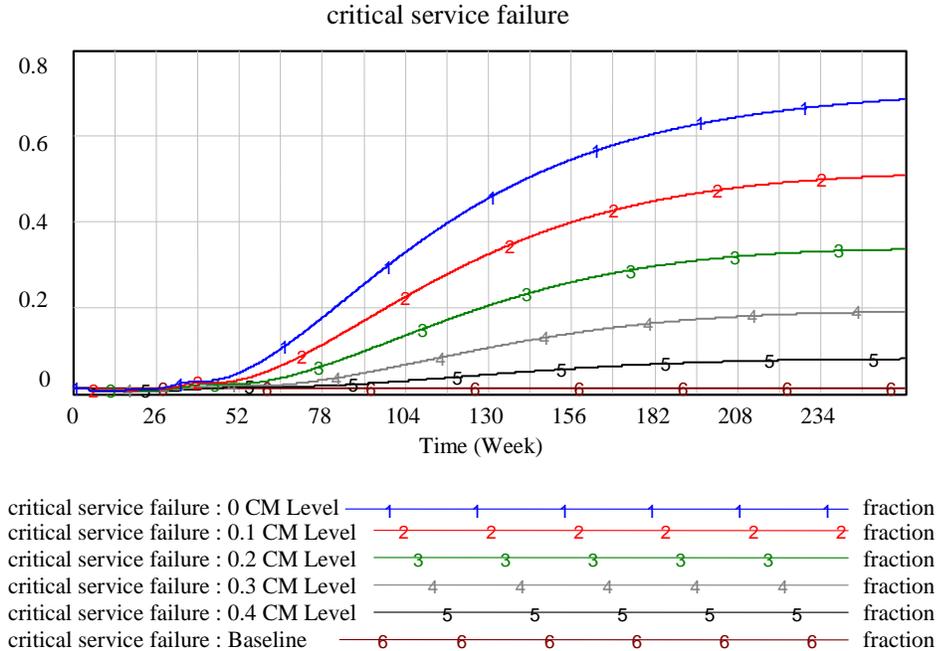
Figure 16: Testing Levels of Change Control Lower than Normal

Figure 17 shows the simulation results with levels of change control higher than normal. We expect that the higher change controls would result in lower critical service failure in the long term. This appears to be the case for levels between 0.5 and 0.8. But surprisingly, levels of 0.9 change control and higher result in levels of critical service failure higher than the equilibrium level (which was level 0.5 change control).

critical service failure

critical service failure : Baseline	—1——1——1——1——1——1—	fraction
critical service failure : 0.6 CM Level	—2——2——2——2——2——2—	fraction
critical service failure : 0.7 CM Level	—3——3——3——3——3—	fraction
critical service failure : 0.8 CM Level	—4——4——4——4——4—	fraction
critical service failure : 0.9 CM Level	—5——5——5——5——5—	fraction
critical service failure : 1 CM Level	—6——6——6——6——6——6—	fraction
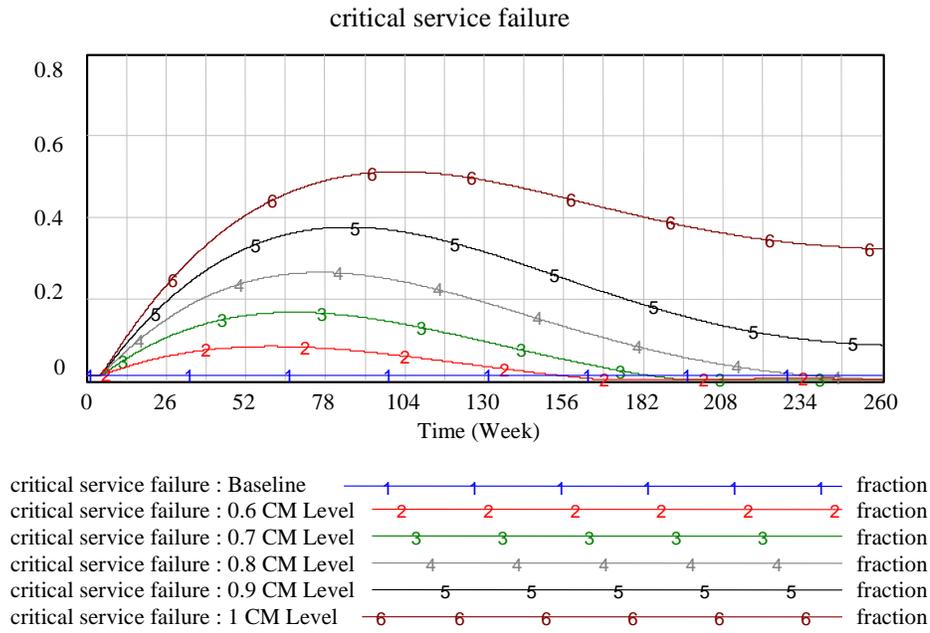
Figure 17:   Testing Levels of Change Control Higher than Normal

Figure 18 verifies that model simulation for change control levels between 0.5 and 0.8 does, in fact, achieve lower levels of critical service failure.

critical service failure

critical service failure : Baseline	—1——1——1——1——1——1—	fraction
critical service failure : 0.55 CM Level	—2——2——2——2——2——2—	fraction
critical service failure : 0.6 CM Level	—3——3——3——3——3——3—	fraction
critical service failure : 0.7 CM Level	—4——4——4——4——4——4—	fraction
critical service failure : 0.8 CM Level	—5——5——5——5——5——5—	fraction
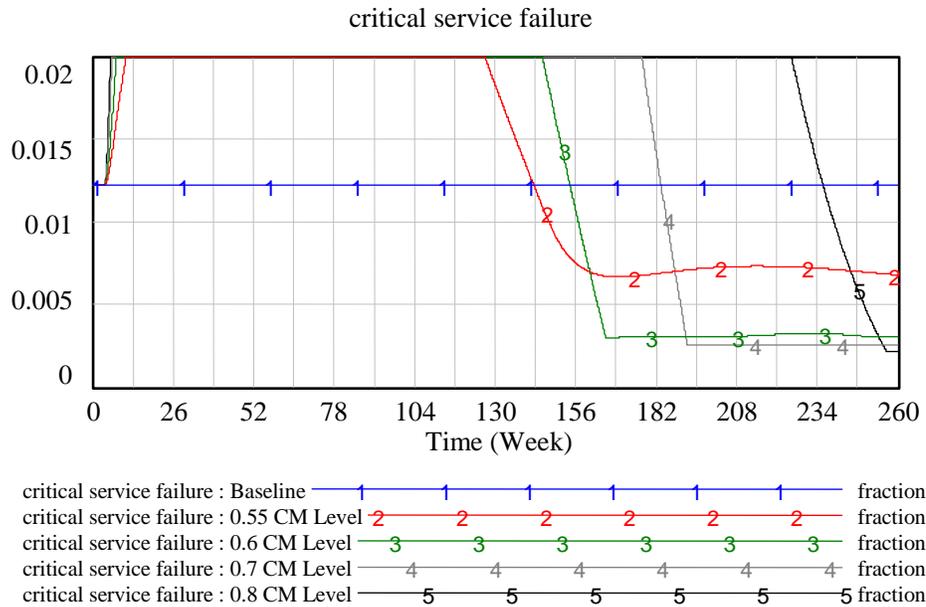
Figure 18: Closer Look for Change Control Between 0.5 and 0.8

Further tests showed that the tipping point between reduced critical service failure and increased critical service failure is a level of change control somewhere between 0.8 and 0.85. Above this level change controls become bureaucratic, that is, excessive change controls cost more than they are worth. One can also see from the above graph the diminishing returns from increased levels of change control. We have yet to determine the optimal level.

The above analysis begs for a characterization of what a certain level of change control actually means in the real world. In future work we hope to use data from the ITPI IT Controls Benchmarking Survey to help with this characterization. That is if we know what constitutes bureaucratic change controls based on the ITPI data, we could characterize the above 0.85 change control level seen above.

4.3 EXTENDED RESULTS WHEN FINDING AND FIXING FRAGILE ARTIFACTS

For the purposes of comparison with previous simulation results we test the model with the same perturbation of its input as above: the step increase in *vulnerabilities discovered* in organizational systems. We test the same combination of organizational responses to policies as before:

- C and nC, depending on whether the organization is committed to its change control policy

- A and nA, depending on whether the organization is committed to its access control policy

- F and nF, depending on whether the organization allows shifting of planned work personnel to problem-repair work

This time, however, we test this model with explicit organizational efforts to find and fix fragile artifacts in place. This will allow comparison with the results described previously where there were no explicit efforts to find and fix fragile artifacts.

Figure 19 shows the *critical service failure* that results over time with a 50% rise in vulnerabilities at the fourth week in the simulation. The baseline run, displayed in blue and labeled 1, shows the system to be in equilibrium with respect to the level of failure. The eight combinations of the above policies are reflected in the eight runs (in addition to the baseline) displayed in the figure.

We make several observations about the simulation runs in Figure 19.

- The use of change and access controls continues to bring about a worse-before-better performance similar to that seen in the case where there was no explicit finding and fixing of fragile artifacts.

- All of the management responses did better, at least in the short term, in the case where the organization made finding and fixing fragile artifacts an explicit part of the planned work.

- Responses that did not permit personnel to be shifted from planned to problem work performed significantly better when organizations explicitly found and fixed fragile artifacts. This is primarily due to the fact that planned-change personnel are the ones finding and fixing fragile artifacts. Therefore, every person taken off of planned work is one less person to find and fix fragile artifacts.

The above suggests that finding and fixing fragile artifacts is an important part of an organization's program to maintain service levels even in the face of external disruptions, such as the 50% increase in exploitable vulnerabilities that we tested.
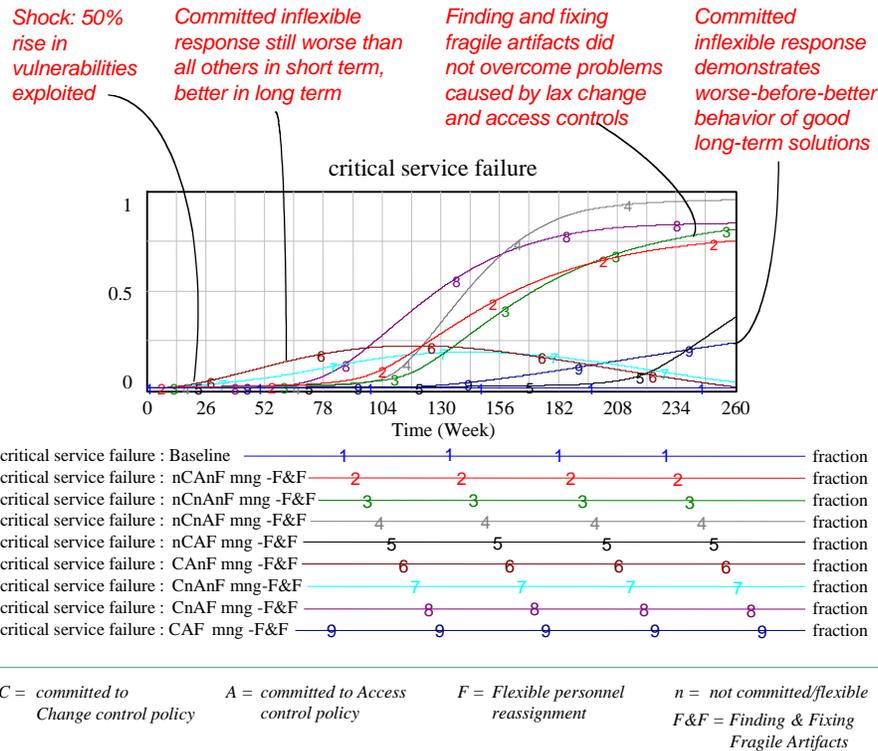
Figure 19: Results from Increasing Vulnerability Discovery by 50% for Critical Service Failure while Finding and Fixing Fragile Artifacts

Figure 20 shows the general benefit of finding and fixing fragile artifacts for the CAnF run. Run 1 shows the results with no concerted efforts to deal with fragile artifacts. Run 2 shows the results when, at week four, the organization starts finding and fixing fragile artifacts as part of their planned work.

The fact that the model has not yet been strongly calibrated based on existing data and expert review suggests that these results might not hold for our final model. However, it does suggest that there may be a need for relaxing particular controls in a regulated way in order to moderate short-term and long-term performance. The benefits of finding and fixing fragile artifacts, however, seem clear, and we expect

the benefits to be substantiated in our continuing modeling efforts, as well as with ongoing applications in the real world.
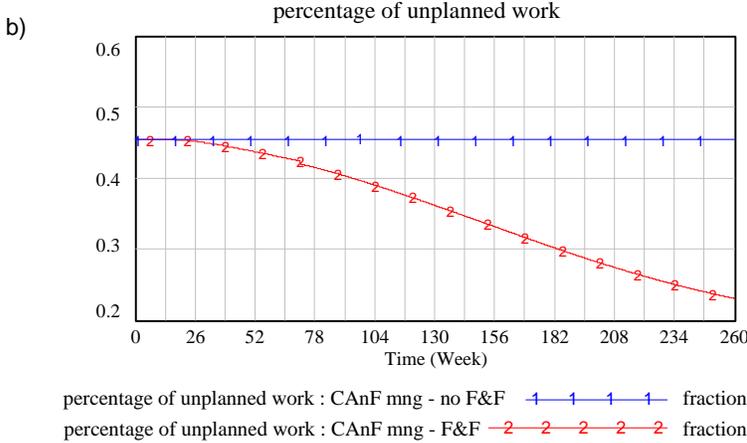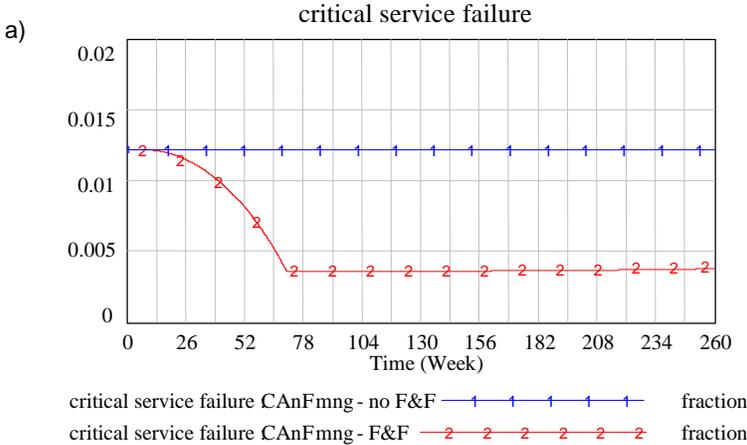
a)

critical service failure

critical service failure CAnFmng - no F&F   1   1   1   1   1   fraction

critical service failure CAnFmng - F&F   2   2   2   2   2   2   fraction

b)

percentage of unplanned work

percentage of unplanned work : CAnF mng - no F&F   1   1   1   1   fraction

percentage of unplanned work : CAnF mng - F&F   2   2   2   2   2   fraction

Key: *C = committed to   A = committed to Access   F = Flexible personnel   n = not committed/flexible*
       Change control policy   control policy   reassignment   F&F = Finding and Fixing
                                                                            Fragile Artifacts

Figure 20: Benefits of Finding and Fixing Fragile Artifacts

5   Conclusions

This report presents an overview of CERT progress in developing a system dynamics model of organizations' typical use of change and access controls to support IT operations. We believe that these models will help organizations understand, specify, and justify a prescriptive process for integrating change and access controls into their business processes in a way that improves security, efficiency, and effectiveness. The execution of these models will help communicate why the foundational controls are effective and provide evidence for the construction of a business case for their adoption.

In summary, we make the following observations associated with our modeling and analysis efforts to date.

- The use of change and access controls is subject to a worse-before-better performance. Some early throughput gains result when these controls are not used, but the long-term advantages of using them outweigh the short-term disadvantages of nonuse. Managers must be aware of the short-term disadvantages so they can persevere through them to accrue the long-term advantages.

- Increasingly rigorous change controls are subject to diminishing returns. Beyond a certain point, change controls become bureaucratic in that their costs outweigh their benefits.

- Shifting personnel from planned work to problem-repair work to manage downtime can work in the short term, but at long-term costs that can overwhelm an organization's ability to successfully manage critical IT service failure. Some discriminate shifting of personnel may be needed to achieve short-term goals, but care must be taken not to sacrifice long-term performance. Future work will test the tradeoffs inherent in this approach.

- Finding and fixing fragile artifacts is an effective way to improve performance regardless of whether other IT controls are used.

- Responses that do not permit personnel to be shifted from planned to problem work bring significantly better performance when organizations explicitly find and fix fragile artifacts.

- Difficulties associated with assessing the fragility of organizational systems and with reducing that fragility suggest that a program of finding and fixing fragile artifacts is best performed in combination with the use of IT controls.

The improved performance through the use of IT controls that is demonstrated by the model simulation assumes that an organization has limited resources to put into problem management.

5.1 DISCUSSION

The problematic behavior patterns that we have described in this paper are similar to the behaviors specified in Repenning and Sterman's paper on problems with sustaining process improvement within organizations [Repenning 2001]. Repenning and Sterman convincingly argue that process improvement efforts have low success rates in organizations not because of any inherent deficiency in the techniques themselves, but because of "how the introduction of a new improvement effort interacts with the physical, economic, social and psychological structures in which implementation takes place." They show that workers shortcut (often covertly) process improvement attempts when work pressure runs high to keep pace with production demands. Wiik makes similar observations with regard to improving the effectiveness of computer security incident response teams [Wiik 2005].

Whether the shortcut is scrimping on a new process improvement technique or, as in our case, on change and access controls already in place in the organization, the result is the same: near-term performance improves and long-term performance declines. In our case, the shortcuts involve reduced change testing and documentation and relaxed staff access controls on operational system artifacts. These shortcuts work to improve system availability in the short term by expediting the problem repair process. This improvement reinforces workers' belief that their shortcuts are helpful thus increasing the likelihood that they'll take the same actions when the next crisis hits. It also makes it difficult for the workers to go back to the more rigorous controls after the immediate crisis is over.

Unfortunately, as we have seen, our model indicates that shortcuts on IT change and access controls are subject to better-before-worse performance. System performance declines only after a significant time has elapsed following the imposition of the shortcut. But people generally assume that cause and effect are closely related in time [Forrester 1994]. So workers and managers often miss the connection between the shortcuts taken and the worsening performance. In addition, business managers often over-emphasize worker deficiencies as the cause for problems and under-emphasize the environmental influences. This tendency, known in attribution theory as the *fundamental attribution error*, means that managers will often associate problems with worker personality shortcomings such as laziness rather than the need to provide sufficient time to allow workers to adhere to a disciplined work process. As a result managers put even more pressure on workers to produce and workers take even more shortcuts because they believe them to be effective. This creates a self-reinforcing spiral toward lower and lower performance (or more and more heroic effort needed to maintain a certain level of performance) in the long term.

The above suggests that most low-performing organizations will have a difficult time adopting and sustaining IT change and access controls without significant efforts to educate IT personnel on (1) the extent to which they sacrifice long-term benefits when they scrimp on these controls and (2) the psychological, social, and

economic forces that act on them as they try to adopt and sustain rigorous change and access controls.

5.2 FUTURE WORK

Our future work will focus both on model refinement and confidence building. Two questions must be answered based on review of the current model:

1. Are we modeling the right things?

2. Are we asking the right questions of the model?

Confidence building is needed to make sure that we have confidence in what the model is telling us. Three questions are also important here:

1. Are the parameters to the model accurate?

2. Are the relationships between components of the model accurate?

3. Is the performance over time predicted by model simulation reasonable and justifiable?

Clearly, efforts to improve confidence in the model may require model refinement. The appropriate mix of model refinement and confidence-building effort will depend on feedback from our sponsors and other readers of this report

We view this report as a checkpoint for our current progress and future plans. Feedback on this report is important to ensure that we are following a path consistent with the overall efforts. Ultimately, we expect that providing the IT management and audit communities with these models and simulations will provide a fact-based approach to determining which controls are foundational, catalytic, and contribute most to simultaneously reducing security risk and increasing effectiveness and efficiency. This work will help create the foundational basis and the first principles that could be useful towards creating guiding principles for IT operational excellence.

Appendix A: Model Assumptions

Our model should characterize an organization that is representative of the class (or a subclass) of organizations (low performers) that we are trying to influence. We are interested in characteristics of those organizations that are important for the domain of application of change and access controls in IT management. The following outlines the primary assumptions that we have made thus far. They are subject to change based on feedback and model refinement.

Organizational Staffing

1. Staff includes IT development staff and IT operations staff.

2. IT development staff includes planned-change personnel.

3. IT operations staff includes problem repair personnel.

4. No new personnel are hired—personnel may only shift between the two staffs. Note: While this may not be particularly realistic, organizations cannot always hire more people to help solve their IT problems. The point of our current model is to see how well an organization can do with staff on hand, where that staff starts out at a reasonable level.

5. Initial state of simulation is as follows:

   – 8.5 problem-repair personnel (on average in one week each person can fix 35 artifacts and diagnose one service failure).

   – 3 planned-change personnel (on average one person can upgrade a service in 25 weeks).

Services

6. Services may be in a state of operation, a state of upgrade, or a state of failure/repair.

7. Unplanned work involves (emergency) IT problem repair of a failed (or degraded) service due to

   – vulnerability exploitation

   – usage stress

   – malfunctioning hardware or software

8. Planned work involves (non-emergency) IT planned changes to upgrade a service for

   – business service extension

   – business service modification/evolution

   – non-emergency vulnerability repair (e.g., the failure of one server of a redundant pair, where the service keeps running despite the failure)

9. Initial state of simulation is as follows:

    – ~52 critical services: 50 operational, 1 being upgraded, and .5 failed. (On average 2% of the operational services fail every week; every 25 weeks a service is upgraded.)

    – % of unplanned work = 45%

    – user standard service failure = 1%. (This is the percentage of services currently accepted by the users.)

Artifacts:

10. Vulnerabilities in artifacts include technical vulnerabilities and hardware faults.

11. Vulnerabilities in artifacts create unreliable artifacts that are either resolved through planned changes or unplanned work arising from service failures.

12. Unreliable artifacts cause service failures.

13. The higher the number of unreliable artifacts per service the more often they fail.

14. Problem repair of a service involves artifact fixes of unreliable artifacts.

15. Service upgrade involves planned changes.

16. *Failure diagnosis rate* is distinguished from *failure repair rate*. Likewise, *unreliable artifact discovery* is distinguished from *artifact fix rate*.

    – Increasing the level of documentation increases diagnosis and discovery rates, but decreases rates of failure repair and artifact fix (because they need to be documented).

    – Increasing the level of testing increases the success of planned changes and artifact fixes, but decreases rates of planned changes and failure repair (including artifact fix rate).

    – Increasing the level of access control decreases spurious artifact corruption, but also decreases the rate of planned changes and failure repair.

    – Increased levels of testing result in increased levels of documentation.

    – Access controls increase the overhead of testing.

17. The overall system documentation quality depends on the quality of planned change documentation and the quality of artifact fix (problem repair) documentation. Poor planned change documentation can compound the problems caused by poor fix documentation.

18. Unauthorized changes due to lack of access controls among the development and operations staff may conflict with one another and lead to a multiplicative effect on artifact corruption.

19. Initial state of simulation is as follows:

    – 156,000 artifacts (3K/service): 10% unreliable (i.e., 300 unreliable artifacts/service), 1% fragile (i.e., 30 fragile artifacts/service)

- % of change success = 72% (50% artifact fix success; 90% planned change success)
- At normal levels of access control on the development and operations staff (initially set at 50%), an average of 25 artifacts/week (out of 156K) are corrupted, that is, made unreliable.

Fragile artifacts

20. The presence of fragile artifacts increases the chance that changes will fail, all other things being equal.

21. Fragile artifacts are introduced only as a result of planned and unplanned changes.
    - Conflicts between failure repairs and service upgrades lead to a multiplicative effect on fragile artifact introduction.
    - 1 artifact is introduced every 2 weeks at a service upgrade rate of 0.12 services/week and a failure repair rate of 1 service per week.

22. A certain (low) volume of finding fragile artifacts occurs as a natural result of operational problems.
    - 0.5% of fragile artifacts are discovered per week without explicit attempts at discovery.

23. A certain (low) volume of fixing fragile artifacts occurs as a natural result of the service upgrade process.
    - 1 artifact is fixed every 2 weeks at a service upgrade rate of 0.12 services/week.

24. Explicit attempts to find and fix fragile artifacts result in larger volumes of artifacts made nonfragile.
    - 2.5% of fragile artifacts are discovered per week with explicit attempts at discovery.
    - 1 artifact is fixed per week at a service upgrade rate of 0.12 services/week.

25. Initial state of simulation is as follows:
    - 0.25% fragility (390 fragile artifacts)
    - 100 fragile artifacts not discovered
    - 290 fragile artifacts discovered

Theory

26. Work pressure on the IT operations and development staff may cause shifting of personnel and reductions in change and access control.

Appendix B: Complete Systems Dynamics Model of Change and Access Controls
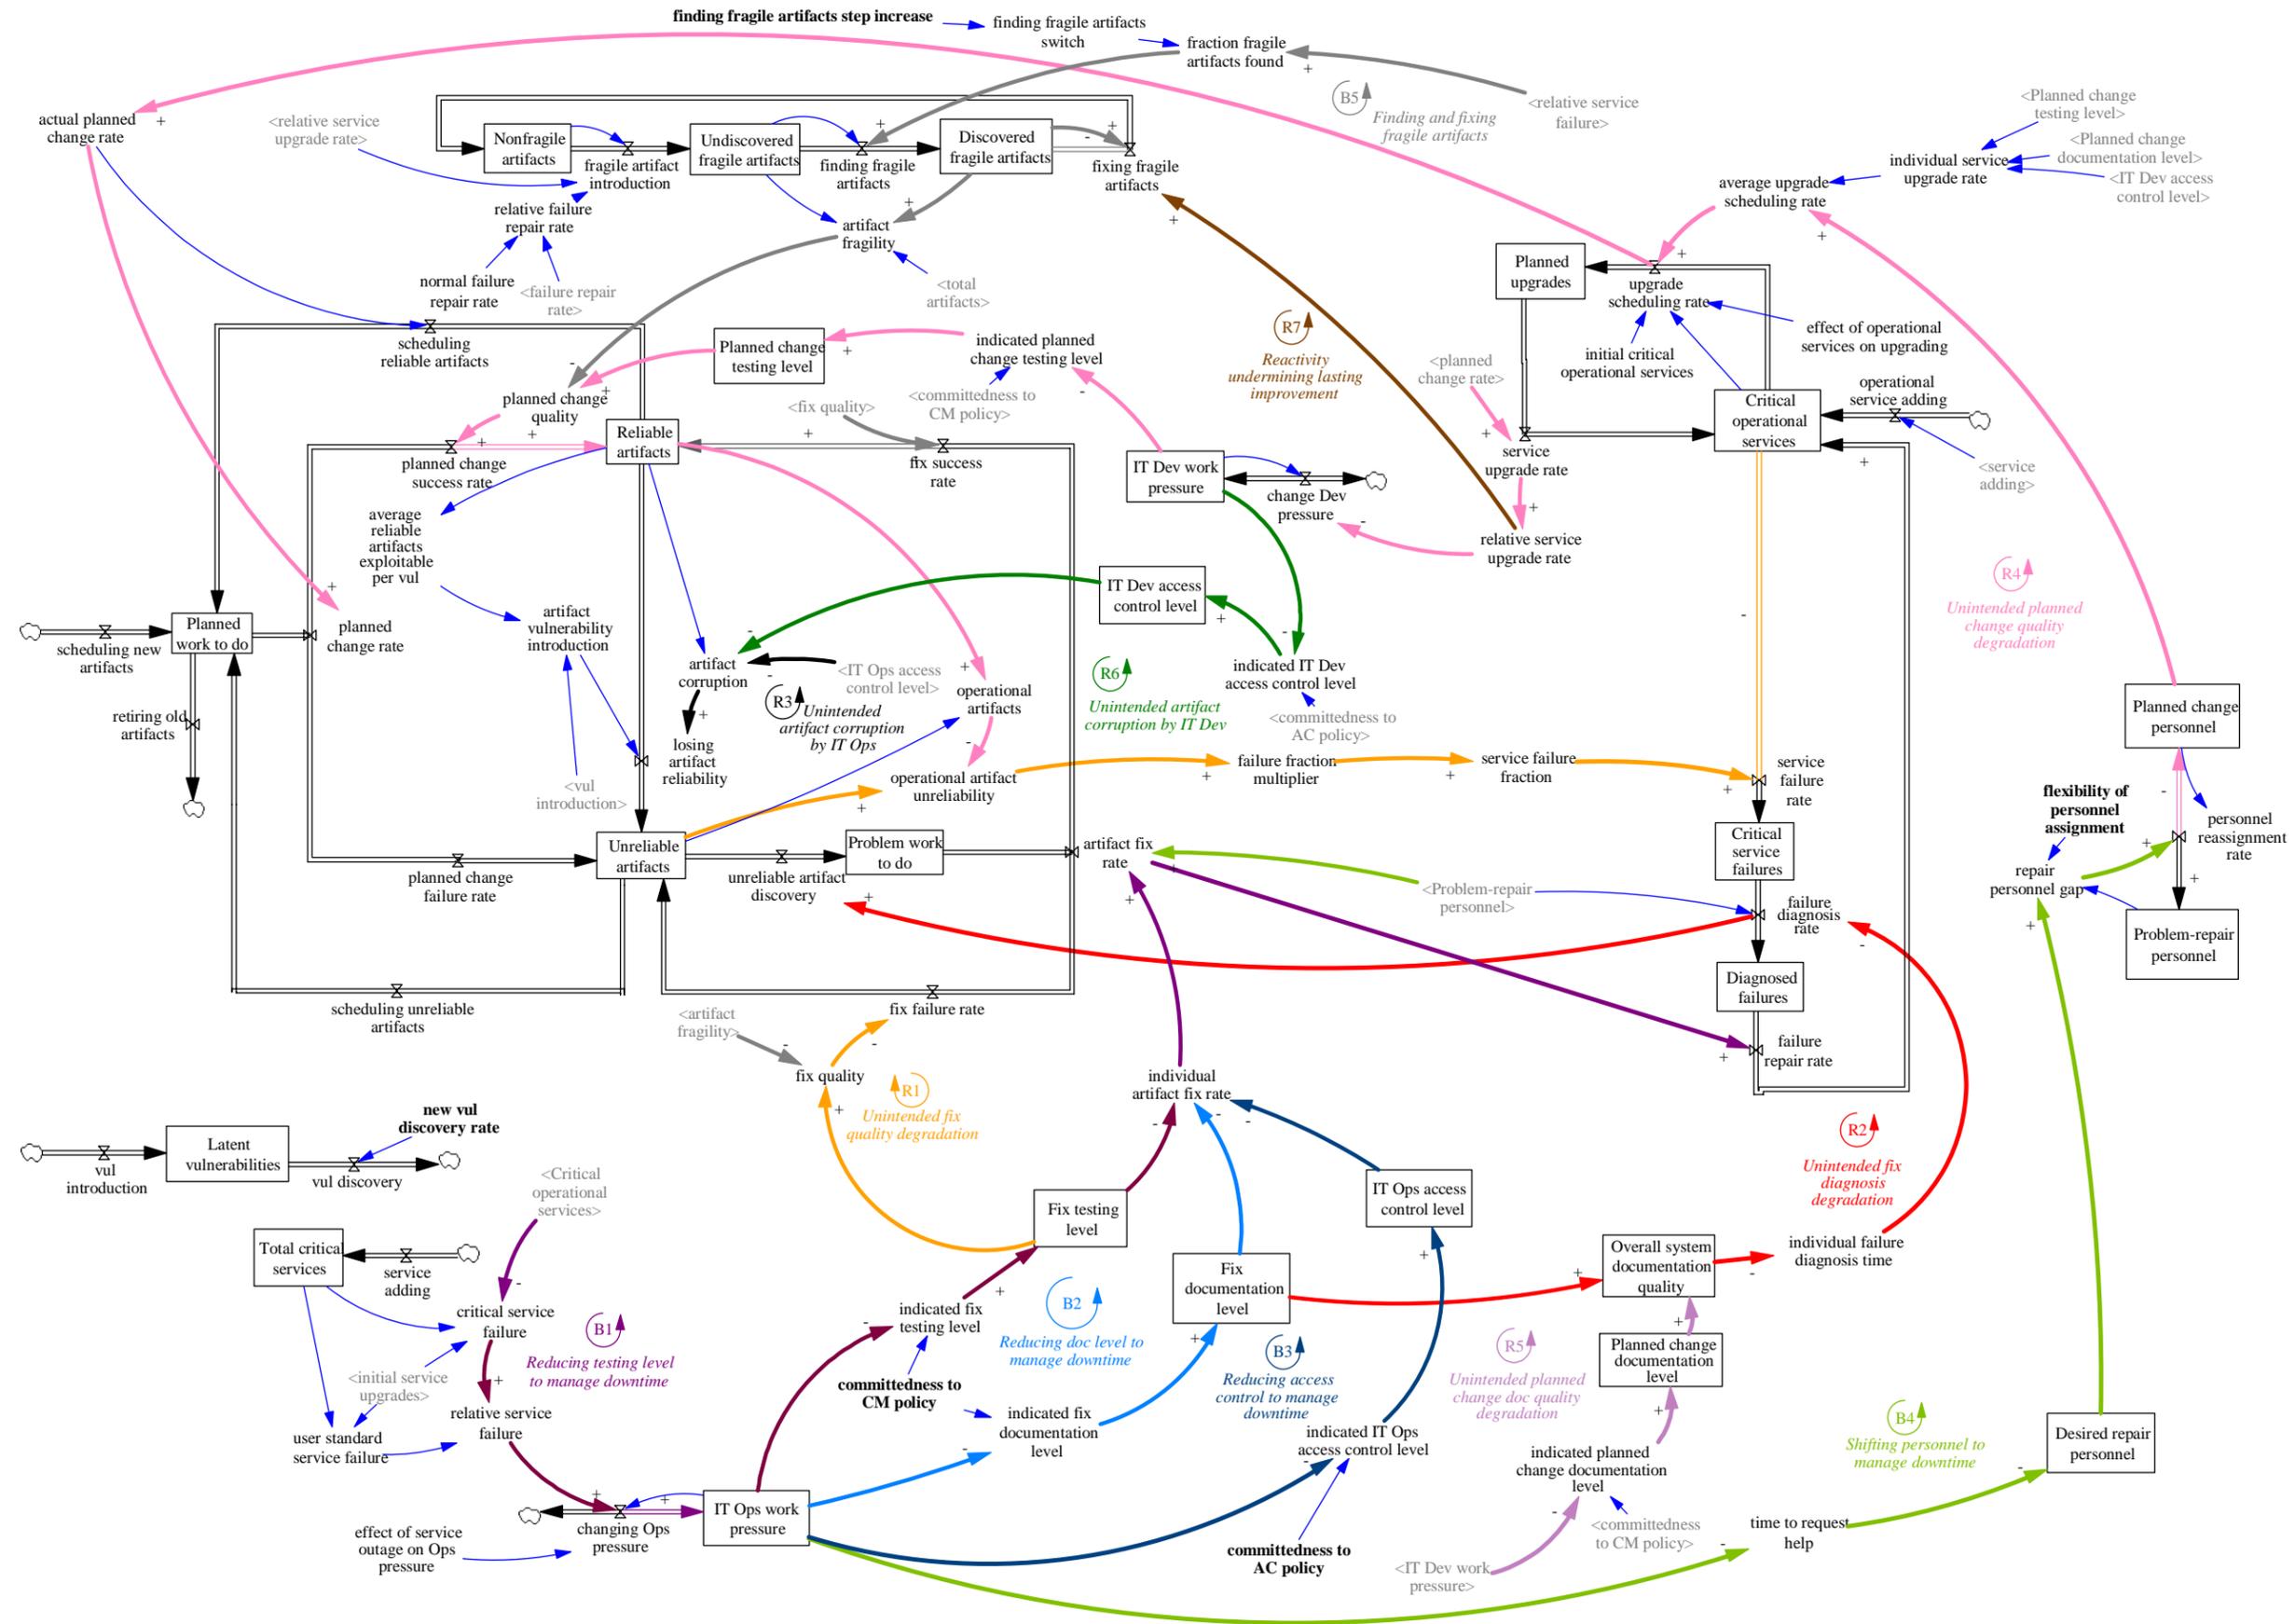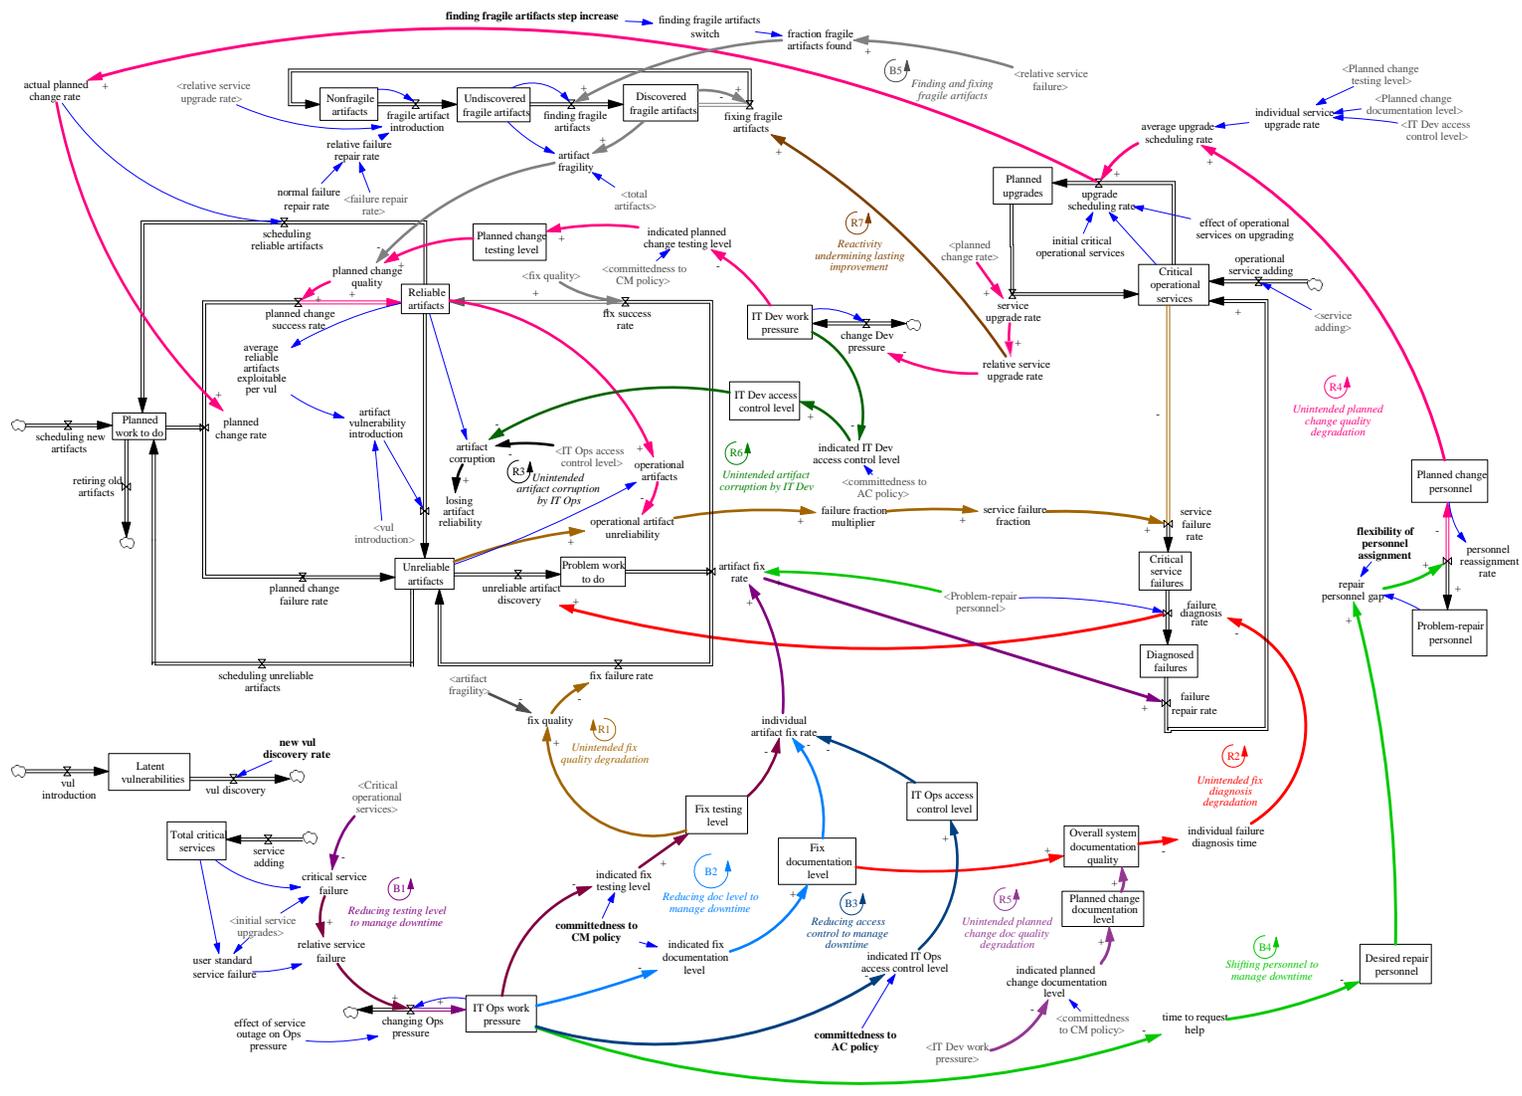
References

[Antao 2005]
Antao, R. S. "Performance Improvement through Change and Access Control Integration." PhD diss., Information Networking Institute, Carnegie Mellon University, 2005. http://www.cert.org/archive/pdf/PICA060119.pdf.

[Barabba 2002]
Barabba, Vince; Huber, Chet; Cooke, Fred; Pudar, Nick; Smith, Jim; & Paich, Mark. "A Multimethod Approach for Creating New Business Models: The General Motors OnStar Project," *Interfaces 32*, 1 (January 2002): 20-34.
http://interface.highwire.org/cgi/content/abstract/32/1/20.

[Behr 2005]
Behr, K. N.; Castner, G.; & Kim, G. "Quantifying the Value, Effectiveness, Efficiency, and Security of IT Controls." Information Technology Process Institute, 2005.
http://www.itpi.org/docs/ITPI_Controls_Benchmarking_Survey_Initial_Findings_v0817.pdf.

[Behr 2004]
Behr, K.; Kim, G.; & Spafford, G. *Visible Ops Handbook: Starting ITIL in Four Practical Steps*. Information Technology Process Institute, 2004.

[Brenner 2002]
Brenner, Michael; Radisic, Igor; & Schollmeyer, Maria. "A Criteria Catalog Based Methodology for Analyzing Service Management Processes," 145-156. *Management Technologies for E-Commerce and E-Business Applications: 13th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, (DSOM 2002)*. Montreal, Canada, Oct. 2002. Berlin, Germany: Springer, 2002.
http://www.springerlink.com/content/x5gry7uh5m4tllj3/.

[Forrester 1994]
Forrester, Jay W. "Learning through System Dynamics as Preparation for the 21st Century." Keynote address. June 27-29, 1994. Systems Thinking and Dynamic Modeling Conference for K-12 Education, Concord Academy, Concord, MA.
http://sysdyn.clexchange.org/sdep/papers/D-4434-3.pdf.

[Garvin 1995]
Garvin, D. A. "The Process of Organization and Management." *MIT Sloan Management Review 39*, 4 (Summer 1998): 33-50.

[IIA 2004]
The Institute of Internal Auditors (IIA). *Information Technology Controls.* Altamonte Springs, FL: The Institute of Internal Auditors.
http://www.theiia.org/download.cfm?file=70284.

[ITPI 2004a]
IT Process Institute (ITPI). *ITPI Controls Benchmarking Survey*.
http://www.itpi.org/home/veesc.php (2004).

[ITPI 2004b]
IT Process Institute (ITPI). *The VEESC survey of practice: A call to action*.
http://www.tcpipservices.com/Default.aspx?tabid=55 (2004).

[ITPI 2005]
IT Process Institute (ITPI). *Uncover the Business Value of IT Controls.*
http://www.itpi.org/docs/The%20IT%20Process%20Institute%20Benchmarking%2
0Survey%20Brochure%20v6.pdf (2005).

[ITPI 2007]
IT Process Institute (ITPI). "The IT Process Institute: Advancing the Science of IT
Management." http://www.itpi.org/home/aboutus.php (2007).

[Kim 2005]
Kim, G. & Warmack, R. "Proving Control of the Infrastructure." *Sarbanes-Oxley
Compliance Journal* (July 2006). http://www.atvcapital.com/news.php?id=249.

[Meadows 1974]
Meadows, Dennis L.; Behrens, William W.; Meadows, Donella H.; Naill, Roger F.;
Randers, Jergen; & Zahn, Erich K. O. *Dynamics of Growth in a Finite World*,
Cambridge, MA: Productivity Press Inc., 1974 (ISBN 0-262-13142-0).

[Repenning 2001]
Repenning, Nelson P. & Sterman, John D. "Nobody Ever Gets Credit for Fixing
Problems that Never Happened: Creating and Sustaining Process Improvement."
California Management Review 43, 4 (Summer 2001): 64-88.

[Sarvann 2000]
Sarvanan, Devaraj & Kohli, Rajiv. *The IT Payoff*: *Measuring the Business Value of
Information Technology Investment.* Indianapolis, IN: FT Press, 2002 (ISBN 0-
130-65074-9).

[Senge 1990]
Senge, Peter M. *The Fifth Discipline: The Art & Practice of The Learning Organi-
zation*. New York, NY: Random House, 2006 (ISBN 0-385-51725-4).

[SoftLanding 2005]
SoftLanding Systems. *The Sarbanes-Oxley Act FAQ.*
http://slseurope.com/sox/docs/sox-faq.pdf (2005).

[Sterman 2000]
Sterman, John D. *Business Dynamics: Systems Thinking and Modeling for a Com-
plex World*. Columbus, OH: McGraw-Hill, 2000 (ISBN 0-072-38915-X).

[Stern 2001]
Stern, Andrea. "Reinvesting the IT Dollar: From IT Firefighting to Quality Strate-
gic Services." *Educause Quarterly 24,* 3 (2001): 8-14.

[Taylor 2005]
Taylor, Jay R; Allen, Julia A.; Hyatt, Glenn L.; & Kim, Gene H. *Change and Patch Management Controls: Critical for Organizational Success.* Altamonte Springs, FL: The Institute of Internal Auditors, 2005 (ISBN 0-894-13574-0*)*.

[Wiik 2005]
Wiik, Johannes; Gonzalez, Jose J.; & Kossakowski, Klaus-Peter. "Limits to effectiveness of Computer Security Incident Response Teams (CSIRTs)." In Proc*eedings of the 23rd International Conference of the System Dynamics Society*. Boston, MA: The System Dynamics
Society.

[Wolstenholme 2003]
Wolstenholme, Eric F. "Towards the Definition and Use of a Core Set of Archetypal Structures in System Dynamics." System Dynamics Review *19*, 1 (Spring 2003): 7–26.

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE March 2007	3. REPORT TYPE AND DATES COVERED Final

4. TITLE AND SUBTITLE Modeling and Analysis of Information Technology Change and Access Controls in the Business Context	5. FUNDING NUMBERS FA8721-05-C-0003

6. AUTHOR(S)

Andrew P. Moore & Rohit S. Antao

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2006-TN-040

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116	10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS	12B DISTRIBUTION CODE

13. ABSTRACT (MAXIMUM 200 WORDS)

Ongoing field work centered at the Information Technology Process Institute (ITPI) makes clear that processes that control change and access within information technology (IT) management and operations simultaneously reduce security risk and increase efficiency and effectiveness. The CERT® Coordination Center is building on this work. This technical note describes a system dynamics model that embodies CERT's current hypothesis of why and how these controls reduce the problematic behavior of the low-performing IT operation. CERT has also started to extend the model in ways that reflect the improved performance seen by high performers. In the longer term, the hope is that this model will help to specify, explain, and justify a prescriptive process for integrating change and access controls into organizations' business processes in a way that most effectively reduces security risk and increases IT operational effectiveness and efficiency.

14. SUBJECT TERMS Change control, access control, information technology, IT, system dynamics model, IT management	15. NUMBER OF PAGES 59

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL