Improving Management of Information Technology: System Dynamics Analysis of IT Controls in Context

Andrew P. Moore
CERT, Software Engineering Institute
Carnegie Mellon University
4500 Fifth Avenue
Pittsburgh, PA 15213
Tel: (412) 268-5465
apm@cert.org

Rohit S. Antao
734 S. 4th Street,
Philadelphia, PA, 19147
Tel: (412) 708-2746
rohit.antao@us.pwc.com

Abstract

Ongoing field work centered at the Information Technology Process Institute (ITPI) is finding that change and access wheel simultaneously reduce security risk and increase the efficiency and effectiveness of information technology (IT) management and operations.[1] *The CERT*[®] *Coordination Center is building on this work. This paper describes a system dynamics model that embodies our current hypothesis of why and how these controls reduce the problematic behavior of the low-performing IT operation. We have also started to extend the model in ways that reflect the improved performance seen by high performers. In the longer term, we hope this model will help to specify, explain, and justify a prescriptive process for integrating change and access controls into their 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 builds on foundational work done at the Information Technology Process Institute (ITPI) by investigating the causal relationship between IT performance and change and access controls (Behr 2005). 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

[®] The CERT Coordination Center is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

[1] *Controls* are processes that provide assurance for information and information services, and help mitigate risks associated with technology use. *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).

with stringent regulatory requirements (Castner 2005). IT, 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 is 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.[2] In addition to coping with demands of effectiveness and efficiency from businesses, IT must satisfy regulatory requirements issued by laws such as the United States Sarbanes Oxley Act of 2002. 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 assigned 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 (Allen 2004)**.**

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).

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. Base on these experiences ITPI hypothesizes that not all IT controls contributed equally to IT effectiveness, efficiency, and security. Those that do we call foundational controls in the sense that 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 (ITPI 2004). This 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

[2] *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).

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).

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	Satisfactory security	3	Low amounts of unplanned work
		4	High change rates
		5	High change success rates
		6	Low number of repeat audit findings
	Efficiency		**Efficiency**
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	% of IT budget spent on compliance
		4	% of IT budget spent on operations

Table 1: High performance indicators

1.3 Our Research

Motivated by the conflicting positions on the efficacy of change and access controls in IT performance, our research attempts 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 (Moore 2005). This behavior arises from an inability, or even purposeful negligence, to take into consideration the long-term effects or unanticipated consequences of the decision to bypass these 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 overhead 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 motivate. 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.

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.

Table 2: Expected benefits of this research

2 Representative Case: MediTAir, Inc.

This section describes MediTAir, Inc., a fictional airline experiencing growing pains as it tries to expand its IT infrastructure to support an increasing client base. Using a fictional case as a basis for the modeling activity allows us to characterize the prototypical low performer and assumptions on which the model is based in a compelling way without referring to the performance of a real organization. It also characterizes the class of problematic operations that could be improved through the use of more rigorous change and access controls. Any resemblance to a real organization or sector is purely accidental. We conclude this section by hypothesizing the cause of MediTAir’s performance problems.

2.1 Company Overview

Founded in April 2000, with its headquarters in New York, MediTAir is a regional airline that operates in the Mediterranean. With a fleet of 20 50-seat CRJ-200s, it operates approximately 35 flights daily between Cyprus and Crete, Crete and Corsica, France and Italy, and Spain and Madeira, as well as a weekly service between Cyprus and Cairo.

In addition to passengers, it also carries freight and express packages on its passenger flights and has interlined small cargo freight agreements with various other local carriers. The company also offers its customers package deals, allowing them to make hotel bookings as well as arrange for airport pickups and car rentals.

MediTAir markets its services to its customers and interacts with its business partners primarily through the internet. Management continues to publicize this as a competitive advantage. A few sales, however, are also made through travel agencies and the airline's reservation call centers.

In December 2001, MediTAir went public and over the past year has been trying to expand its operations as well as improve its quality and range of services through its online portal.

MediTAir's Market and Competitive Edge

Although MediTAir's primary market includes American tourists traveling to the Mediterranean region, the company's marketing team has seen a growing popularity among European and Australian tourists traveling in the region.

MediTAir has been able to keep its ticket prices low because it is a small airline with a minimal ground staff. In addition, the vast majority of its customers' reservations are made online.

Company Culture

MediTAir employs a dynamic and talented group of people, some of whom have been with the company since its inception, and others who were hired as it grew. Recruitment focuses on well-educated technical and business professionals with reputations for high performance. The company fostered a culture of "do whatever it takes" to get the job done. Initiatives taken by employees were highly rewarded and all employees were subject to quarterly performance reviews that were tied directly to their compensation. Unsuccessful managers did not last long and those who made "heroic" efforts were highly regarded.

In its early years, MediTAir managed compliance-related risk by relying on legal counsel and undergoing periodic regulatory audits. However, faced with increasing governmental regulations and penalties for failing to comply, management's culture of compliance is gaining new meaning. Corporate governance, risk, and comprehensive corporate compliance are moving to the forefront of the corporate boardroom dialogue. Compliance is now viewed as an enterprise-wide function. The company wants to use technology to implement the best possible practices to efficiently manage compliance, thereby reducing both the cost and the risk of falling out of compliance. Over the past two years, MediTAir has spent $5 million on compliance, and has hired a Chief Compliance Officer as well as a team of 5 internal auditors and another 10 consultants.

Business Demands on IT Operations

With the growing popularity of MediTAir among European and Australian tourists, management identified a need to enhance the company's infrastructure to support the additional customer load in November 2002.

The IT department was entrusted with several responsibilities to support this decision, including delivering key new applications to maintain and increase MediTAir's competitive edge, as well as implementing the necessary infrastructure to support the expected increase in the capacity of its reservation system.

Additionally, in an effort to improve the quality as well as the range of services offered, the development team made enhancements to the logistical decision support system and the new customer flight reservation system. IT Operations has been notified about the need to provide infrastructural support regarding these additions.

Current Situation

In January 2004, IT Operations officially began working on these projects; however, the department has been frequently sidetracked by planned[3] and unplanned[4] operational and maintenance issues. By spending most of its time on these issues, the department faces cost and schedule overruns. These costs and costs incurred from downtime during peak seasons have forced management to raise ticket prices. Management also observed a trend: IT Operations was failing to deliver new services to help the business achieve its objectives on time and within budget. Additionally, findings reported by internal audit exposed a number of weaknesses in the IT infrastructure for financial reporting. Management began exploring the idea of outsourcing the work of the entire division. However, it was finally decided to hire an independent consultant to analyze the situation and make recommendations for getting MediTAir back on track.

2.2 Problematic Behavior

In August 2003, MediTAir invested in an expensive Change Management System (CMS). The company observed that changes were being made unchecked, that many changes were failing, and that it was impossible to trace changes back to their source. More importantly, implementation of the system would get the auditors off the company's back for a while.

The company's targeted service level was 99% availability; however, due to an outage that began on Nov. 18 and spanned almost two days, the company achieved only 94% availability for the quarter. Additionally, this outage generated expensive overtime. The outage, it was eventually discovered, resulted from a developer's decision to upgrade 50% of the Web servers with new code for the hotel booking system, changing 93 critical executables. This caused the hotel booking session to freeze the servers. The upgraded servers locked up so hard that they could not even reboot. This caused the site to go down, taking the whole line of business with it.

IT Operations was immediately thrown into firefighting mode—at the time it had no insight into the cause of the problem. The department tried random remedies, none of which worked. In an emergency meeting of everyone from Operations and R&D, no one admitted to making any changes to the system. While the business was down during the peak period, Operations and R&D spent time finger-pointing and mud-slinging. No one took accountability for any change, and without any kind of documentation or detective measures in place, everyone preceded based on suspicions from previous outages.

Eventually, cooler heads prevailed and some useful clues were found. They discovered that a development upgrade had been tentatively explored. Because developers also had access to production servers, the person responsible for the change could have been anyone on the core R&D team. After several phone calls, a possible responsible party was identified, but he had left for the night. At 3 a.m., the development manager contacted the developer, who admitted that he had done a "small upgrade," but insisted that this change could not have caused the outage. When Operations installed the supposedly harmless executables onto a test box, they were

[3] Planned activities range from scheduled upgrades to patch deployment.

[4] Unplanned activities include work generated from failed changes.

immediately able to replicate the problem. It had taken them 30 hours to diagnose the cause of the outage!

MediTAir had more than six nearly catastrophic failures that year, similar to this outage. The most recent problem involved someone from the IT Operations team applying a service pack to the SQL server. This led to the disabling of the billing service and a corresponding five hours of downtime. Too much time was spent on crisis management and, as a result, the department fell behind on planned projects. However, IT Operations personnel did admit that not all of these failures were a result of changes made by the R&D team. In fact, only 70% of the changes made by Operations worked the first time without generating a firefighting episode. Not only were they spending more on unplanned work than planned work, the majority of the unplanned work was self-generated. When services went down or possible vulnerabilities were identified, the department made no effort to prioritize tasks.

Toward the middle of 2004, IT Operations returned to reactive mode as it tried to get its infrastructure to meet the demands of the audit team. A server had been found on the network that had been compromised for two years. Ad hoc, chaotic, and urgent changes were being made to the system to fix problems as they were identified. Audit insisted that Operations should be handling all the production systems, but Operations simply couldn't find the time for it and delegated this responsibility to the developers. Operations began to dread any interaction with the security or audit teams as they only seemed to create obstacles and limitations to getting work done; almost every attempt to satisfy their demands in one section of the network led to breaks in others.

Operations admitted that as pressure increased to satisfy IT Security, IT Audit, and the business mission, documentation of changes decreased. Operations circumvented the CMS to get more "real" work done. Management kept finding excuses to not use this system.

The absence of a structured and systematic way to manage change only worsened the situation. "Do whatever it takes" efforts by members of the Operations and R&D teams helped them achieve their immediate objectives, but came back to bite them later on. With a lack of documentation, they found themselves moving in a downward spiral. In doing so much firefighting, Operations could not find the time for planned activities, such as maintenance and implementing infrastructural support for new services.

After analyzing MediTAir's problematic behavior over time, the reference mode, illustrated in Figure 1, was plotted. Upon studying this pattern they observed that as service disruption increased, the emergency repair work done succeeded in fixing only some of the 'broken' services. This success, however, was short-lived, and as time went by MediTAir began facing even more service disruption than it started out with. As the service disruption faced by the organization gradually increased, emergency or unplanned work went up too. With resources (and the focus) being shifted to emergency work, the fraction of planned project work began to decrease. Another observation made was that the problem resolution time began to increase as time went by.

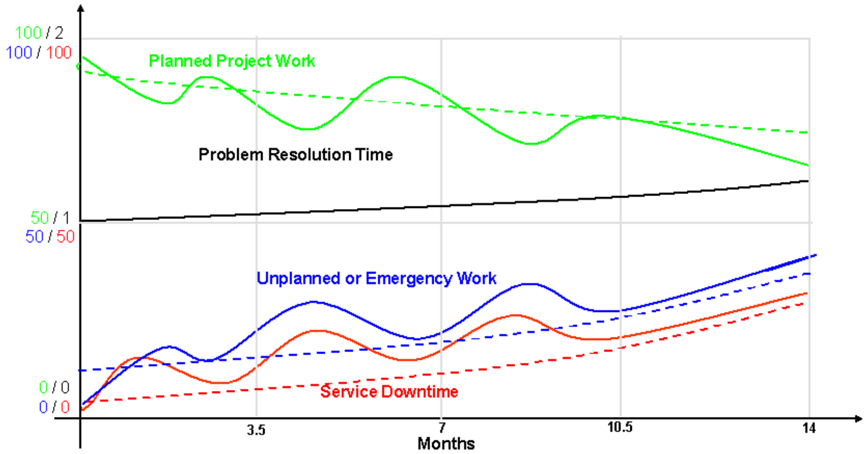

Figure 1: Problematic Behavior Over Time

2.3 Dynamic Hypothesis

We hypothesize that a root cause for the problematic behavior experienced by MediTAir rests in an organizational inclination to deal with system disruption by making decisions to

- relax the enforcement of change and access controls
- shift excessive resources from proactive to reactive work

This tendency stems from an inability, or even purposeful negligence, to take into consideration the long-term effects or unintended consequences of those decisions. There was a disproportionate focus on short-term benefits, as opposed to long-term improvements, leading to an unrelenting cycle of fixes that failed over time. IT managers underestimated the negative impact of the decisions they made because those decisions were separated in time as well as space (Senge 1990).

The result of this (IT) operational approach was a patchwork of official and undocumented workarounds, supporting a patchwork of increasingly unstable and undocumented software and systems that continued to degrade over time. Moreover, the combination of fragmented processes and IT systems conspired to undermine the organization’s ability to understand and control the operational environment, leading to a downward spiral of ever increasing operational problems.

An uncommitted and patchwork approach towards the implementations of these controls made them ineffective, thus preventing organizations from deriving true value from them. This inevitably resulted in these controls being viewed as unnecessary bureaucracy that diminished productivity.

3 System Dynamics Model

MediTAir’s operational environment paints a good picture of the nature of problems faced by many low-performing organizations. We hypothesize that some of the root causes for this low performance were in fact control issues related to ineffective change and access management and an organizational tendency to relocate personnel from proactive to reactive work. To test this hypothesis we need to understand and model the functioning of IT operations within this organization. In this section we attempt to capture the underlying mechanisms of this system as perceived by their operational staff. The model developed here characterizes their mental models; it captures the scenarios under which certain decisions are made and the effects of those decisions, whether intended or not.

The appendix to this paper depicts the full system dynamics model that we have developed. This section provides an overview of that model, which is described more fully in (Moore 2005). 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 the appendix.

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 sequence of linked, interdependent activities, which, taken together, transform inputs to outputs (Garvin 1995).

Take the change management 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 which are documented. In between the input and output phase the requested change progresses through a number of interdependent tasks such as change planning, authorization, testing, documentation, and implementation, as shown in Figure 2 (Behr 2004).

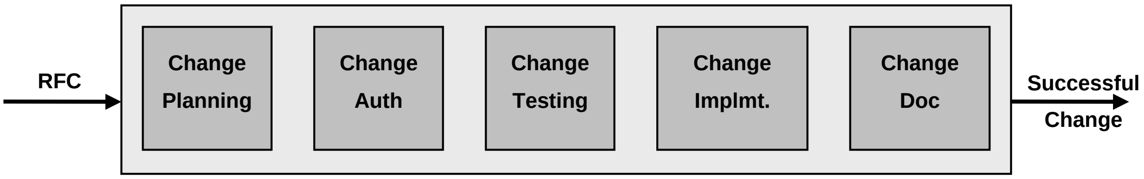

Figure 2: The physical components of the change management process

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.

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

The Service View

Figure 3 shows the service view of the stock and flow model. The *Critical operational services* stock includes those services that are currently up and running. 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[5]. 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

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 but the reason for the failure has not yet been determined. The *Diagnosed failures* stock contains those failed services that have been diagnosed.

The Artifact View

Figure 4 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 with each other. Upgrade scheduling in the service

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

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: (1) vulnerabilities discovered that may be exploited by malicious individual, (2) new unforeseen uses of the artifacts, beyond that for which they were designed, or (3) 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.

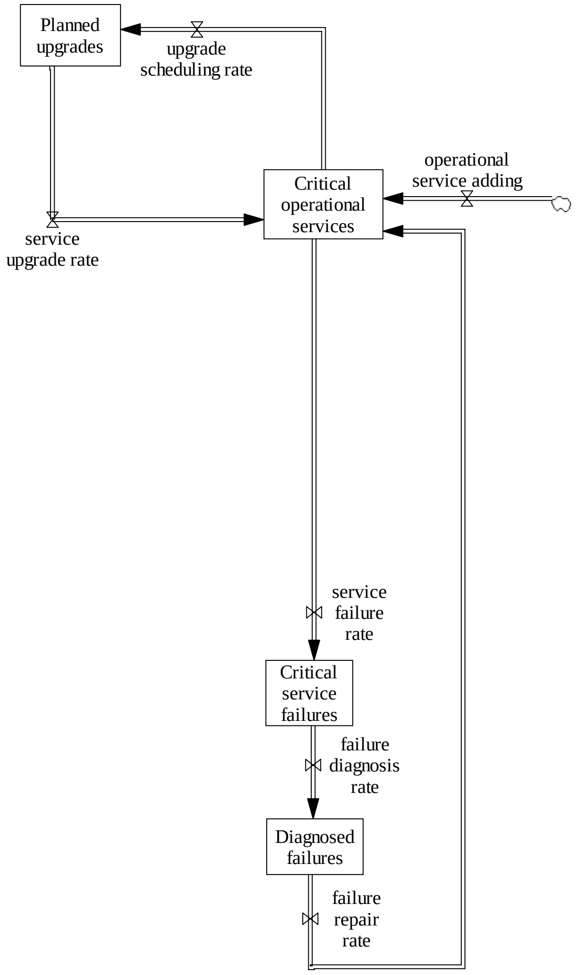

Figure 3: Service Flows

We hypothesize that IT management problems are due to an overly reactive approach that focuses on emergency repair work to keep IT services running. This reactivity 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. Fragile artifacts generate much firefighting and need to be identified and handled with care (Behr 2004).

A high leverage fundamental solution for IT management suffering low performance then should be to find and fix those fragile artifacts that are embedded in their IT infrastructure. Figure 5 depicts an extension to the infrastructure of our system dynamics model. Three stocks of artifacts are added: Nonfragile artifacts, Undiscovered fragile artifacts, and Discovered fragile artifacts. 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 service outage.

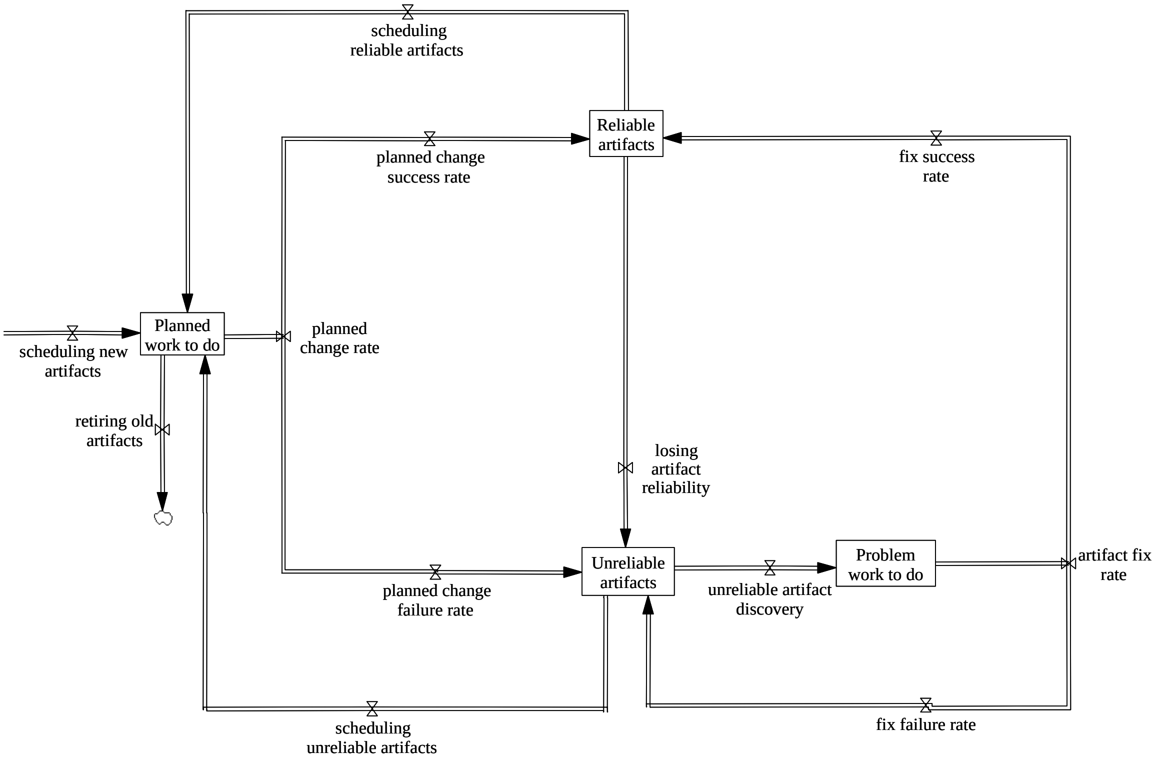

Figure 4: Basic Artifact Flows

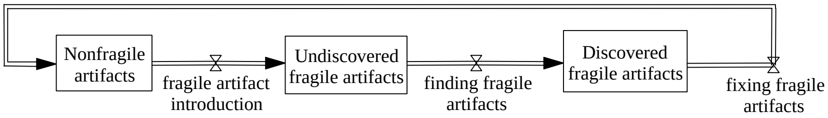

Figure 5: Flows Involving Artifact Fragility[6]

The Personnel View

Figure 6 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.

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

[6] 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 operations, whereas fragile artifacts cause problems due to the stress of change. Of course, an artifact may be both unreliable and fragile.

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.

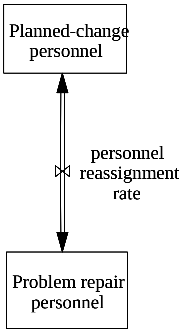

Figure 6: Personnel Flows

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.[7]

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 7, 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. An example of a fix-that-fails in the project-management domain occurs when a project runs behind schedule. A short-term fix might be to force employees to work overtime to get back on schedule—a balancing loop—and an unintended consequence is burn out that can occur due to excessive overtime—a reinforcing loop. Even a fix of hiring more people to get back on schedule may cause unintended consequences in which the experienced people are diverted from their normal "productive" work to train the new hires, thus causing the project schedule to slip even further. While such a fix may work better in the long term, an effective solution may need to be balanced with short-term needs.

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

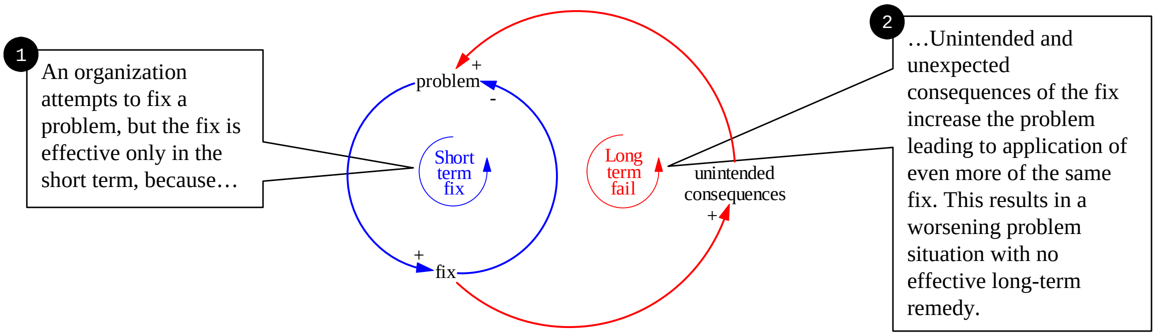

1 An organization attempts to fix a problem, but the fix is effective only in the short term, because…

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 7: 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 8 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 outages and get 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 outages. 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.

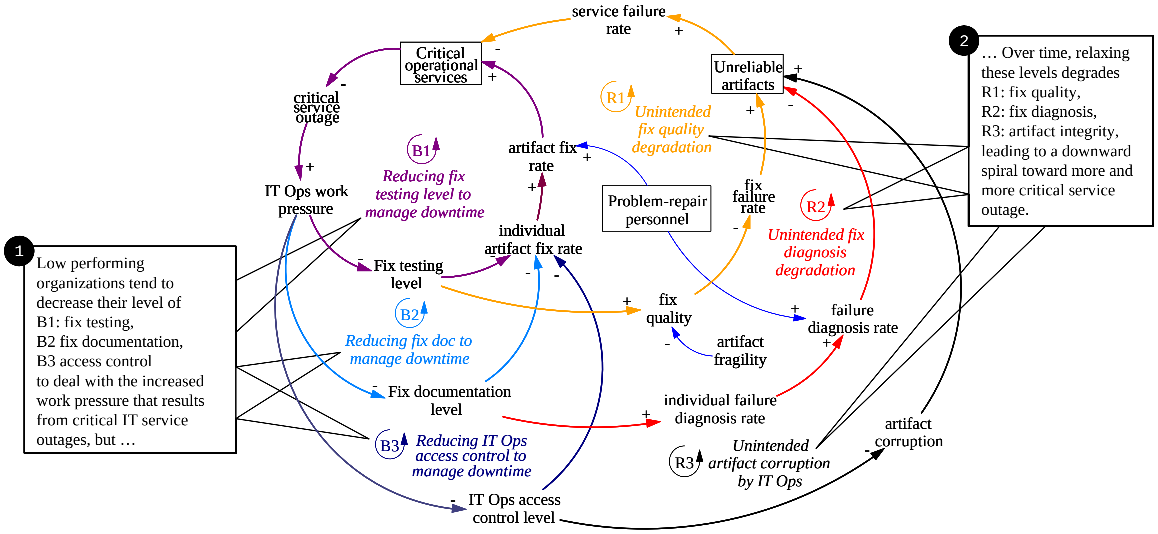
service failure rate
Critical operational services
critical service outage
IT Ops work pressure
B1
Reducing fix testing level to manage downtime
artifact fix rate
individual artifact fix rate
Fix testing level
B2
Reducing fix doc to manage downtime
Fix documentation level
B3
Reducing IT Ops access control to manage downtime
IT Ops access control level
Unreliable artifacts
R1
Unintended fix quality degradation
Problem-repair personnel
fix failure rate
R2
Unintended fix diagnosis degradation
fix quality
artifact fragility
failure diagnosis rate
individual failure diagnosis rate
R3
Unintended artifact corruption by IT Ops
artifact corruption
1
2
+
-

1 Low performing organizations tend to decrease their level of B1: fix testing, B2 fix documentation, B3 access control to deal with the increased work pressure that results from critical IT service outages, but …

2 … Over time, relaxing these levels degrades R1: fix quality, R2: fix diagnosis, R3: artifact integrity, leading to a downward spiral toward more and more critical service outage.

Figure 8: Relaxing Change and Access Controls to Manage Downtime

Figure 9 illustrates the forth 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 to increase failure repair rate and bring 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 unintended consequences, which parallel the unintended consequences seen in Figure 8[8]:

- 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 relaxed, 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

[8] There are also balancing loops in the IT development domain that parallel the “fixes” associated with loops *B1* through *B3* in Figure 8. For simplicity, these balancing loops are not shown.

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.

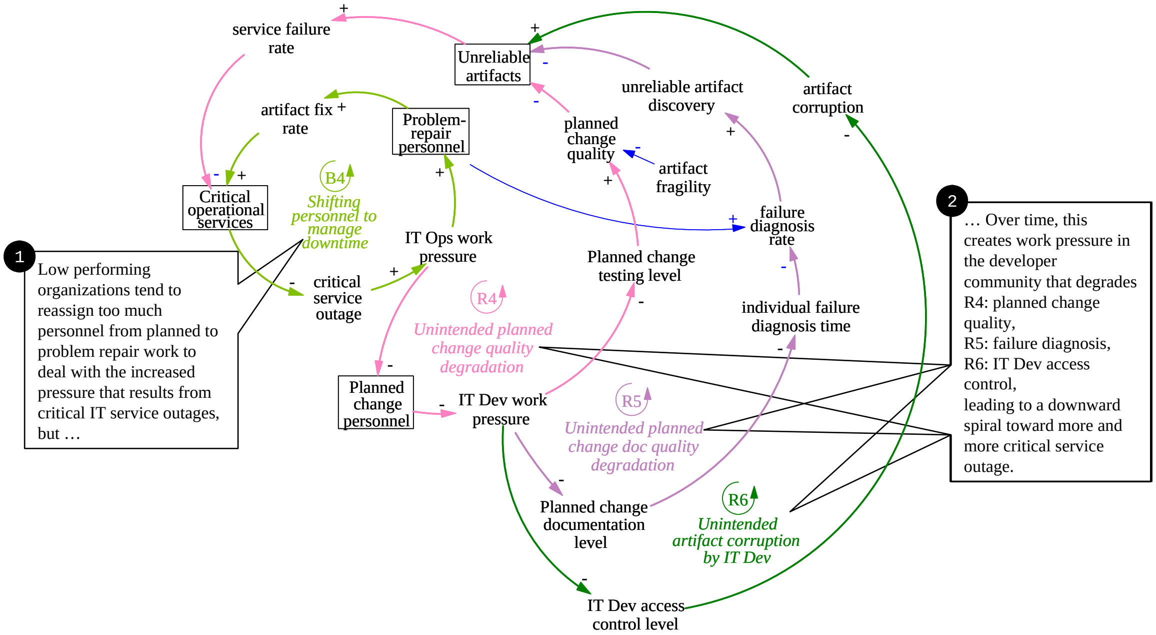
service failure rate
Unreliable artifacts
unreliable artifact discovery
artifact corruption
artifact fix rate
Problem-repair personnel
planned change quality
artifact fragility
Critical operational services
B4
Shifting personnel to manage downtime
IT Ops work pressure
failure diagnosis rate
critical service outage
Planned change testing level
individual failure diagnosis time
R4
Unintended planned change quality degradation
Planned change personnel
IT Dev work pressure
R5
Unintended planned change doc quality degradation
R6
Unintended artifact corruption by IT Dev
Planned change documentation level
IT Dev access control level
1
2
+
-

1 Low performing organizations tend to reassign too much personnel from planned to problem repair work to deal with the increased pressure that results from critical IT service outages, but …

2 … Over time, this creates work pressure in the developer community that degrades R4: planned change quality, R5: failure diagnosis, R6: IT Dev access control, leading to a downward spiral toward more and more critical service outage.

Figure 9: Shifting Planned Change Personnel to Problem Management

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 10 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.

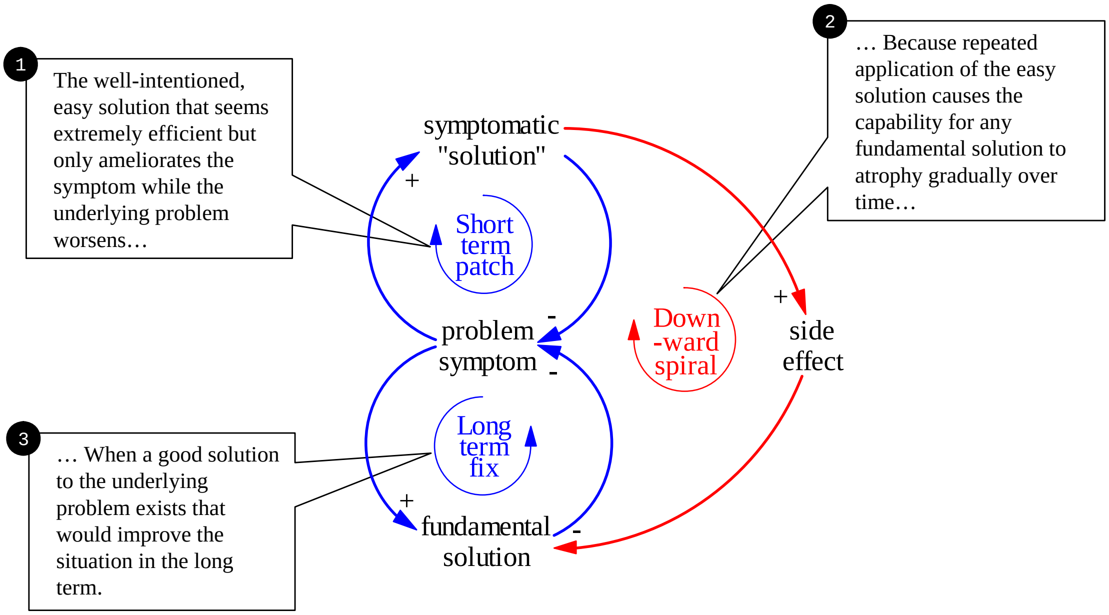
1
The well-intentioned, easy solution that seems extremely efficient but only ameliorates the symptom while the underlying problem worsens…
2
… Because repeated application of the easy solution causes the capability for any fundamental solution to atrophy gradually over time…
3
… When a good solution to the underlying problem exists that would improve the situation in the long term.
symptomatic "solution"
+
Short term patch
problem symptom
-
-
Down -ward spiral
+
side effect
Long term fix
+
fundamental solution
-

Figure 10: Shifting the Burden Archetype

Figure 11 shows two classes of solutions available to IT managers to handle the problem of critical service outage: 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 11 is an example of a symptomatic solution to the problem of service outage. This is the same loop labeled B4 depicted in Figure 9. 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 11 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 change, the change is highly likely to fail. 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 about 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 outage. However, as shown in the *R7a* loop of Figure 11, 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 outage. This worsening situation is characteristic of the Shifting the Burden archetype and the downward spiral of the low performer.

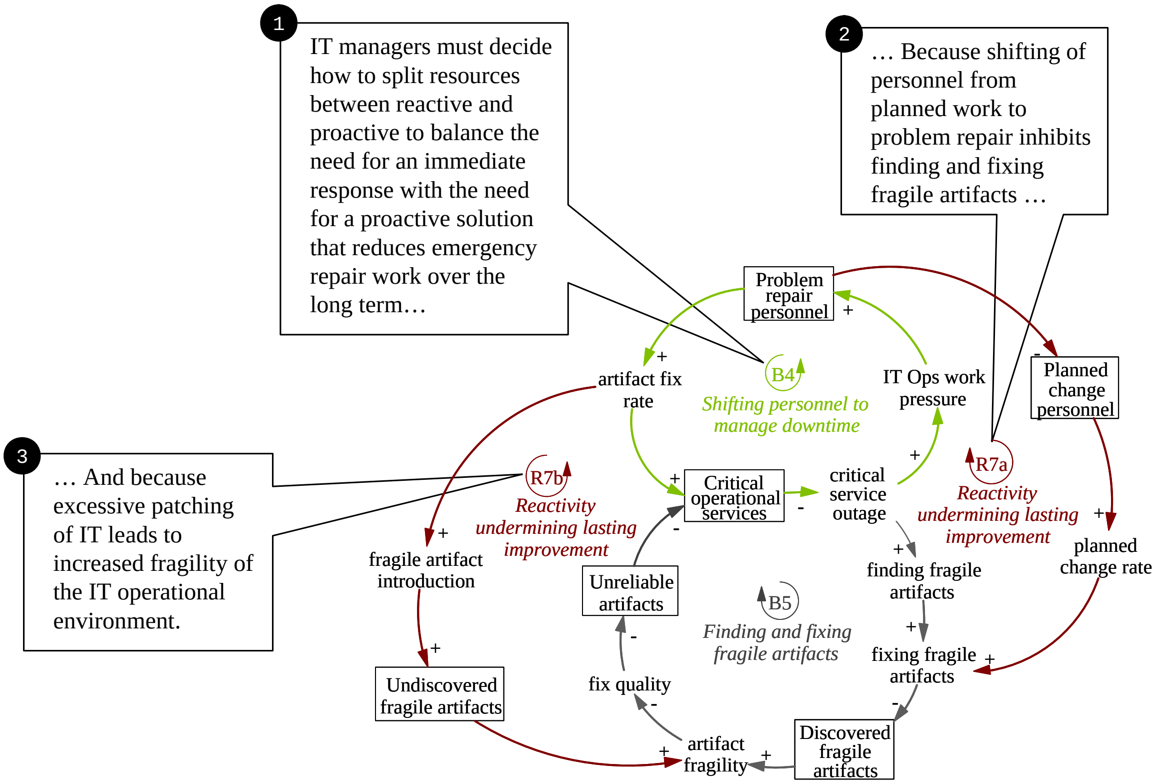
1 IT managers must decide how to split resources between reactive and proactive to balance the need for an immediate response with the need for a proactive solution that reduces emergency repair work over the long term…

2 … Because shifting of personnel from planned work to problem repair inhibits finding and fixing fragile artifacts …

3 … And because excessive patching of IT leads to increased fragility of the IT operational environment.

Problem repair personnel
Planned change personnel
artifact fix rate
B4
Shifting personnel to manage downtime
IT Ops work pressure
R7a
Reactivity undermining lasting improvement
R7b
Reactivity undermining lasting improvement
Critical operational services
critical service outage
planned change rate
fragile artifact introduction
Unreliable artifacts
finding fragile artifacts
B5
Finding and fixing fragile artifacts
fixing fragile artifacts
Undiscovered fragile artifacts
fix quality
artifact fragility
Discovered fragile artifacts
+
-

Figure 11: 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 outage*, intended to be the maximum level of outage users will find generally acceptable.[9]

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, a model is easier to validate and to experiment with since the analyst can more easily 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. Intuitively, this perturbation might arise as a result of 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, staffing of planned work). As work pressures rise, the organization reduces its change controls (respectively, access controls, staffing of planned work) to more quickly 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, staffing of planned work) with the hopes that higher quality fixes and continuance of planned work will pay returns in the long run.

Figure 12 shows the *critical service outage* 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 outage. The rest of the runs, labeled 2 through 9, show the critical service outage 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.

[9] The *user standard service outage* parameter is analogous to a customer-driven requirement for system availability.

We make the following observations about the above runs:

- The use of change and access controls is subject to a worse-before-better behavior. 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.

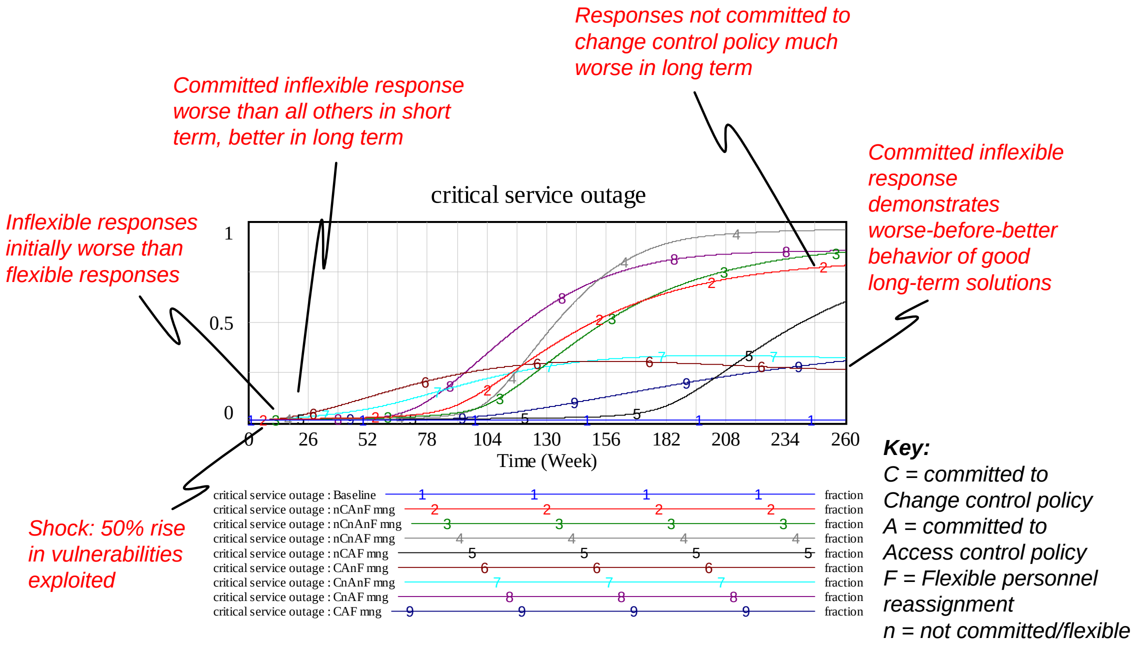

Responses not committed to change control policy much worse in long term

Committed inflexible response worse than all others in short term, better in long term

Committed inflexible response demonstrates worse-before-better behavior of good long-term solutions

Inflexible responses initially worse than flexible responses

Shock: 50% rise in vulnerabilities exploited

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

Figure 12: Results from Increasing Vulnerability Discovery by 50% for critical service outage

Figure 13 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.[10] It is not too surprising that Figure 13a 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.

[10] 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.

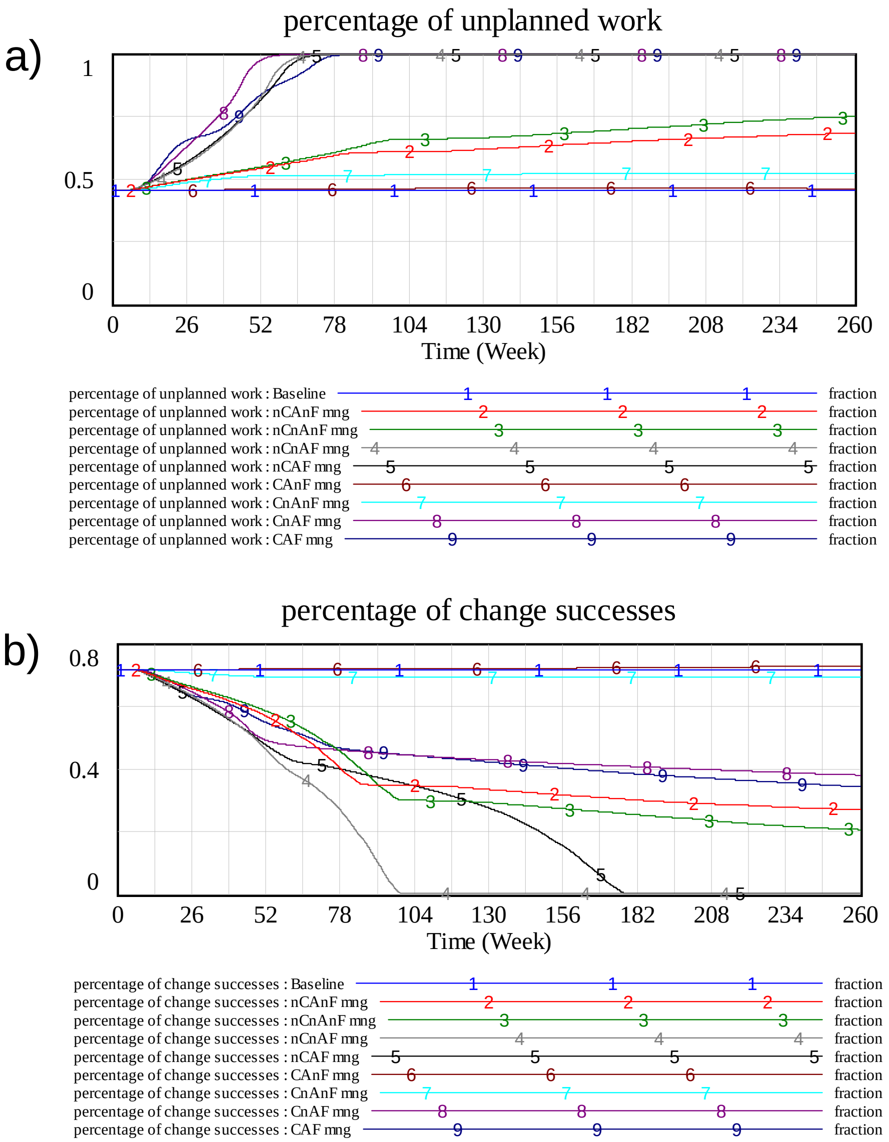

Figure 13: Results from Increasing Vulnerability Discovery by 50% for a) percentage of unplanned work and b) percentage of change successes

The remaining runs of Figure 13a 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 outage that inhibits planned change work.[11] Similarly, the enforcement of 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 13b 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,

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

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 14 verifies that lower levels of change control do lead to greater critical service outage in the model.

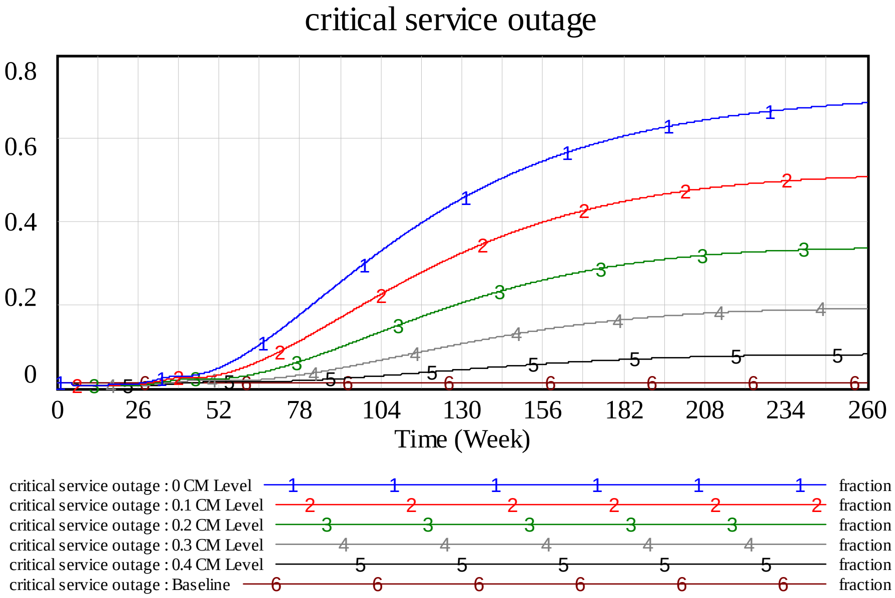

Figure 14: Testing Levels of Change Control Lower than Normal

Figure 15 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 outage 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 outage higher than the equilibrium level (which was level 0.5 change control).

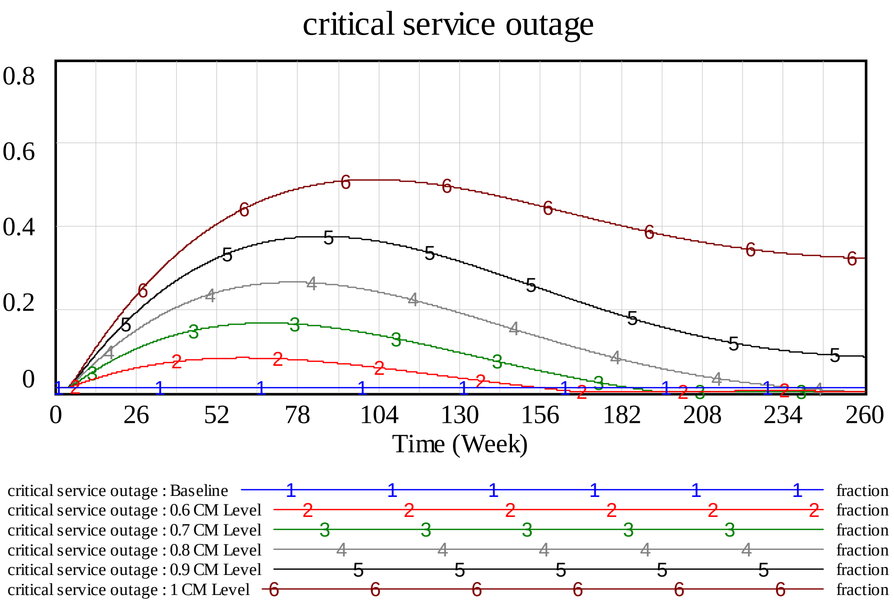

Figure 15: Testing Levels of Change Control Higher than Normal

Figure 16 verifies that model simulation for change control levels between 0.5 and 0.8 does, in fact, achieve lower levels of critical service outage.

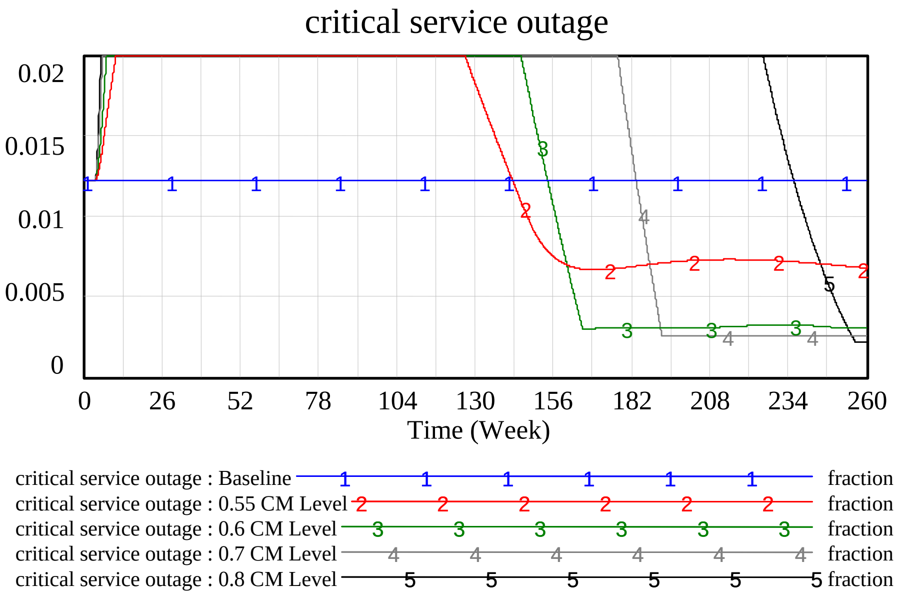

Figure 16: Closer Look for Change Control Between 0.5 and 0.8

Further tests showed that the tipping point between reduced critical service outage and increased critical service outage is a level of change control somewhere between 0.8 and 0.85. Above this level change controls become bureaucratic, i.e., 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 the ITPI survey data to help with this characterization. That is, if we know what bureaucratic change controls based on the survey 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 17 shows the *critical service outage* 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 outage. 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 17.

- The use of change and access controls continues to be subject to a worse-before-better behavior 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 18 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 organizations starts finding and fixing fragile artifacts as part of their planned work.

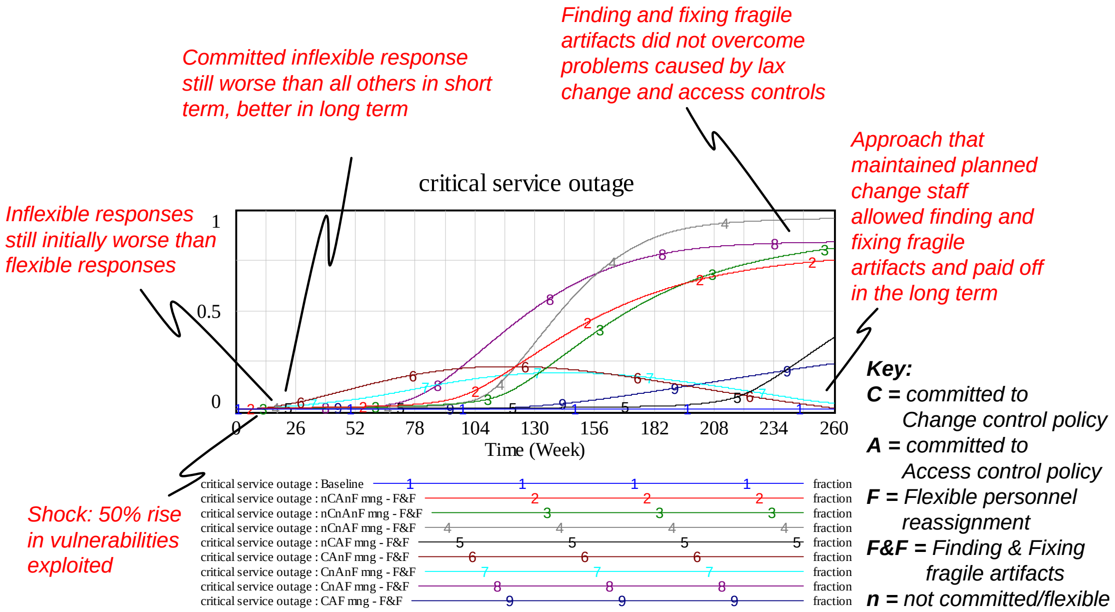

Committed inflexible response still worse than all others in short term, better in long term

Finding and fixing fragile artifacts did not overcome problems caused by lax change and access controls

Approach that maintained planned change staff allowed finding and fixing fragile artifacts and paid off in the long term

Inflexible responses still initially worse than flexible responses

Shock: 50% rise in vulnerabilities exploited

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

Figure 17: Results from Increasing Vulnerability Discovery by 50% for Critical Service Outage while Finding and Fixing Fragile Artifacts

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.

5 Conclusions

This report presents an overview of CERT progress in developing system dynamics models to specify the causal mechanisms underlying observations that change and access controls simultaneously reduce the security risk and increase the efficiency and effectiveness of IT management and operations. Causal models 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. We believe that these models and their execution will help communicate why the foundational controls are effective and provide evidence for the construction of a business case for their adoption.

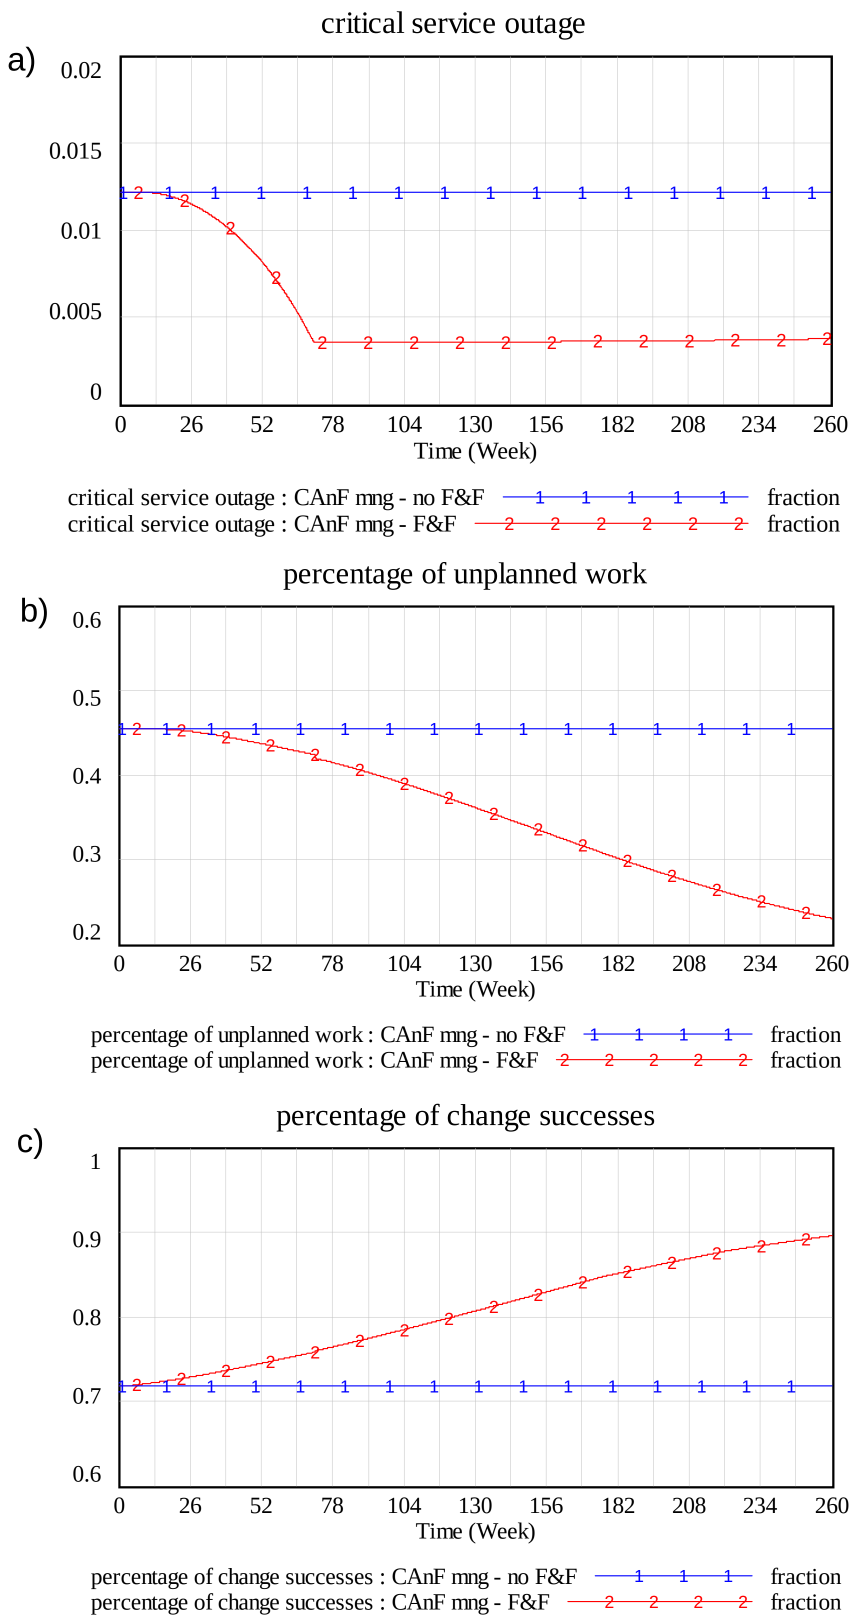

Figure 18: Benefits of Finding and Fixing Fragile Artifacts

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 behavior. Some early throughput gains result when these controls are not used, 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 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 outage. 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 perform significantly better when organizations explicitly find and fix fragile artifacts.
- Difficulties associated with assessing the fragility of organizational systems and with reducing that fragility suggests 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 such 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, change and access controls already in place in the organization, the result is the same: near-term performance improvement and long-term performance decline. 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 control are subject to a better-before-worse behavior. 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, as pointed out by Repenning (Repenning 2001), people 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 them on (1) the extent of the sacrifice that they make 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 focuses both on model refinement and confidence building. Two questions must be answered based on review of the model: 1) Are we modeling the right things? and 2) Are we asking the right questions of the model? Confidence building is needed to make sure that we have confidence in the model results. Three questions are important here: 1) Are the parameters to the model accurate? 2) Are the relationships between components of the model accurate? and 3) Is the behavior 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.

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.

6 Acknowledgements

Many thanks to our ITPI sponsors of this project – especially, Gene Kim and Kurt Milne – and to Julia Allen of the SEI for their support and insights. Also thanks to the conference reviewers of this paper for their comments and pointers to some very useful resources and related work.

7 References

Allen, J., K. Behr, G. Kim et al. 2004. *Best in Class Security and Operations Round Table Report*. Pittsburgh: Software Engineering Institute, Carnegie Mellon University.

Antao, R.S. 2005. "Performance Improvement through Change and Access Control Integration," *Thesis for Master of Science in Information Security Technology and Management, Information Network Institute, CMU,* see http://www.cert.org/archive/pdf/PICA060119.pdf.

Behr, K., G. Castner, G. Kim. 2005. *Quantifying the Value, Effectiveness, Efficiency, and Security of IT Controls*. IT Process Institute. Available from http://www.itpi.org.

Behr, K., G. Kim, and G. Spafford. 2004. *Visible Ops Handbook: Starting ITIL in Four Practical Steps*. IT Process Institute.

Brenner, M., I. Radisic, and M. Schollmeyer. 2002. "A Criteria Catalog based Methodology for Analyzing Service Management Processes". *Montreal: 13th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management*.

Forrester, J.W. 1994. "Learning through System Dynamics as Preparation for the 21st Century," keynote address for *Systems Thinking and Dynamic Modeling Conference for K-12 Education*. Available from http://sysdyn.clexchange.org/sdep/papers/D-4434-3.pdf.

Garvin, D. 1995. *The Process of Organization and Management.* Harvard Business School.

Institute of Internal Auditors. 2004. *Information Technology Controls*. Florida. Available from http://www.theiia.org/index.cfm?doc_id=5166.

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

Kim, G., and R. Warmack. 2005. *Proving Control of the Infrastructure.* Available from http://www.knowledgestorm.com/sol_summary_72899.asp

Meadows, D.L., W.W. Behrens, D.H. Meadows, R.F. Naill, J. Randers, E.K.O. Zahn 1974. *Dynamics of Growth in a Finite World,* Cambridge, MA: Wright-Allen Press, Inc.

Moore, A.P. and R. Antao. 2005. *System Dynamics Modeling and Analysis of IT Management Controls in Context.* Pittsburgh: SEI Special Report CMU/SEI-2005-SR011.

Repenning, N. and J.D. Sterman. 2001. "Nobody Ever Gets Credit for Fixing Problems that Never Happened: Creating and Sustaining Process Improvement," *California Management Review* 43(4): 64-88.

Sarvanan, D., and R. Kohli. 2000. *The IT Payoff*: *Measuring the Business Value of Information Technology Investment.* New Jersey: Prentice Hall.

Senge, P.M. 1990. *The Fifth Discipline*. New York: Doubleday.

Sterman, J.D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. McGraw-Hill.

Stern, A. 2001. "Reinvesting the IT Dollar: From IT Firefighting to Quality Strategic Services," *Educause Quarterly*. 24(3), 8-14.

Taylor, J., J. Allen, G. Hyatt, G. Kim. 2005. *Change and Patch Management Controls: Critical for Organizational Success.* Institute of Internal Auditors.

Wiik, J., J.J. Gonzalez, and K.-P. Kossakowski. 2005. "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, E.F. 2003. "Towards the Definition and Use of a Core Set of Archetypal Structures in System Dynamics," *System Dynamics Review*, 19(1): 7-26.

Appendix: Simulation Model Overview

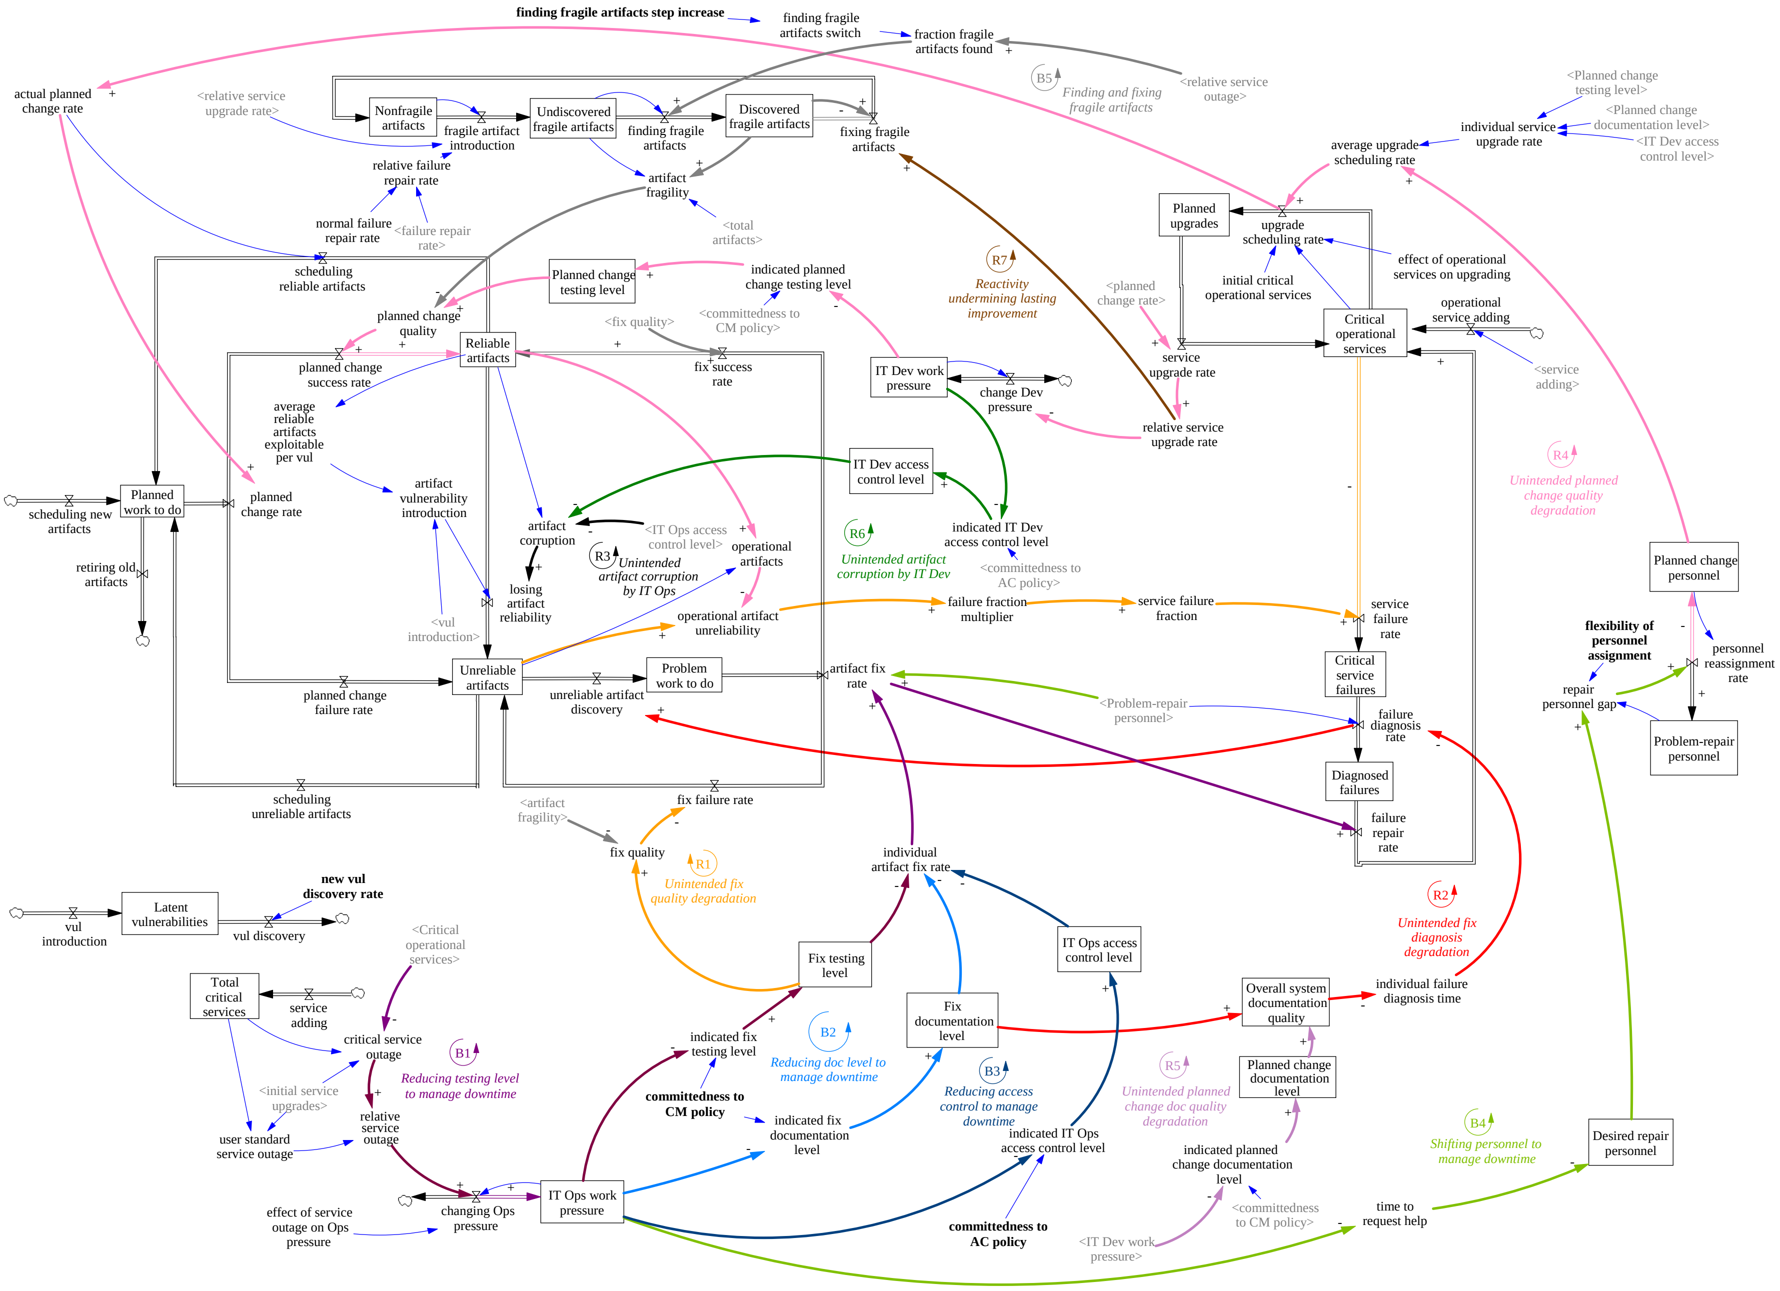
finding fragile artifacts step increase
finding fragile artifacts switch
fraction fragile artifacts found
actual planned change rate
<relative service upgrade rate>
Nonfragile artifacts
fragile artifact introduction
Undiscovered fragile artifacts
finding fragile artifacts
Discovered fragile artifacts
fixing fragile artifacts
B5
Finding and fixing fragile artifacts
<relative service outage>
<Planned change testing level>
<Planned change documentation level>
individual service upgrade rate
<IT Dev access control level>
average upgrade scheduling rate
relative failure repair rate
normal failure repair rate
<failure repair rate>
artifact fragility
<total artifacts>
Planned upgrades
upgrade scheduling rate
initial critical operational services
effect of operational services on upgrading
scheduling reliable artifacts
Planned change testing level
indicated planned change testing level
R7
Reactivity undermining lasting improvement
<planned change rate>
planned change quality
<fix quality>
<committedness to CM policy>
Critical operational services
operational service adding
<service adding>
planned change success rate
Reliable artifacts
fix success rate
IT Dev work pressure
change Dev pressure
service upgrade rate
average reliable artifacts exploitable per vul
relative service upgrade rate
R4
Unintended planned change quality degradation
artifact vulnerability introduction
IT Dev access control level
scheduling new artifacts
Planned work to do
planned change rate
artifact corruption
<IT Ops access control level>
indicated IT Dev access control level
R6
Unintended artifact corruption by IT Dev
R3
Unintended artifact corruption by IT Ops
operational artifacts
retiring old artifacts
losing artifact reliability
<committedness to AC policy>
Planned change personnel
<vul introduction>
operational artifact unreliability
failure fraction multiplier
service failure fraction
service failure rate
flexibility of personnel assignment
personnel reassignment rate
Unreliable artifacts
Problem work to do
artifact fix rate
Critical service failures
repair personnel gap
planned change failure rate
unreliable artifact discovery
<Problem-repair personnel>
failure diagnosis rate
Problem-repair personnel
Diagnosed failures
scheduling unreliable artifacts
<artifact fragility>
fix failure rate
failure repair rate
fix quality
individual artifact fix rate
R1
Unintended fix quality degradation
Latent vulnerabilities
new vul discovery rate
vul introduction
vul discovery
R2
Unintended fix diagnosis degradation
<Critical operational services>
Fix testing level
IT Ops access control level
Total critical services
service adding
Fix documentation level
Overall system documentation quality
individual failure diagnosis time
critical service outage
B1
Reducing testing level to manage downtime
indicated fix testing level
B2
Reducing doc level to manage downtime
B3
Reducing access control to manage downtime
R5
Unintended planned change doc quality degradation
Planned change documentation level
<initial service upgrades>
relative service outage
committedness to CM policy
indicated fix documentation level
indicated IT Ops access control level
B4
Shifting personnel to manage downtime
Desired repair personnel
user standard service outage
indicated planned change documentation level
changing Ops pressure
IT Ops work pressure
<committedness to CM policy>
effect of service outage on Ops pressure
committedness to AC policy
<IT Dev work pressure>
time to request help