veeAM

How to Improve Application Service Delivery and IT Agility

Enhance Availability, application SLAs and utilize your data to improve change management

Contents

01

What You Will Learn

Adapting to Change: The Right Tools for a Proactive ITIL Implementation

Digital transformation offers challenges and opportunities for your business. If you feel outpaced by change, you're certainly not alone.

But with challenges come brand new opportunities. A chance to do things better. Go faster. Be more efficient. If you're treading a tight-rope between innovation and vulnerability, you need to ensure that the choices you make now are not creating risks later. And if you're still firefighting six months down the line, you're doing something wrong.

The Information Technology Infrastructure Library (ITIL) framework offers many opportunities to bring about positive change and improve operational efficiency. It's considered a foundation for best practices in IT infrastructure management because it gives businesses highly effective tools for dealing with aspects like system Availability, IT service continuity, capacity, incident, management and change management.

Welcome to Aperature Labs: a Fictional Case Study

This e-book follows the fictional journey of our ITIL test case, Aperature Labs, struggling with digital transformation, vulnerabilities, and change. Aperature Labs will show us that ITIL works best when it is preventative and proactive, when the right processes are put in place, supported by the right software tools.

A recent study carried out by Forrester found that more organizations (61%) followed ITIL for incident management first and foremost, perhaps ignoring some of its strategic advantages in pre-planning. And Forrester also found less than **a third (29%) of organizations rated themselves as truly proactive in optimizing systems to avoid incidents.**

In this e-book, we will follow Aperature on its journey from having to run ever faster just to stay still, to discovering **how a proactive approach IT management can drive improvements in service and overall business performance**. We will explore how the culture of dealing with IT issues changed, the processes it adopted, and the tools it used.

In this study of a fictional company, Aperature Labs, we'll see how ITIL works best when it is preventative and proactive, when the right processes are put in place supported by the right software tools.
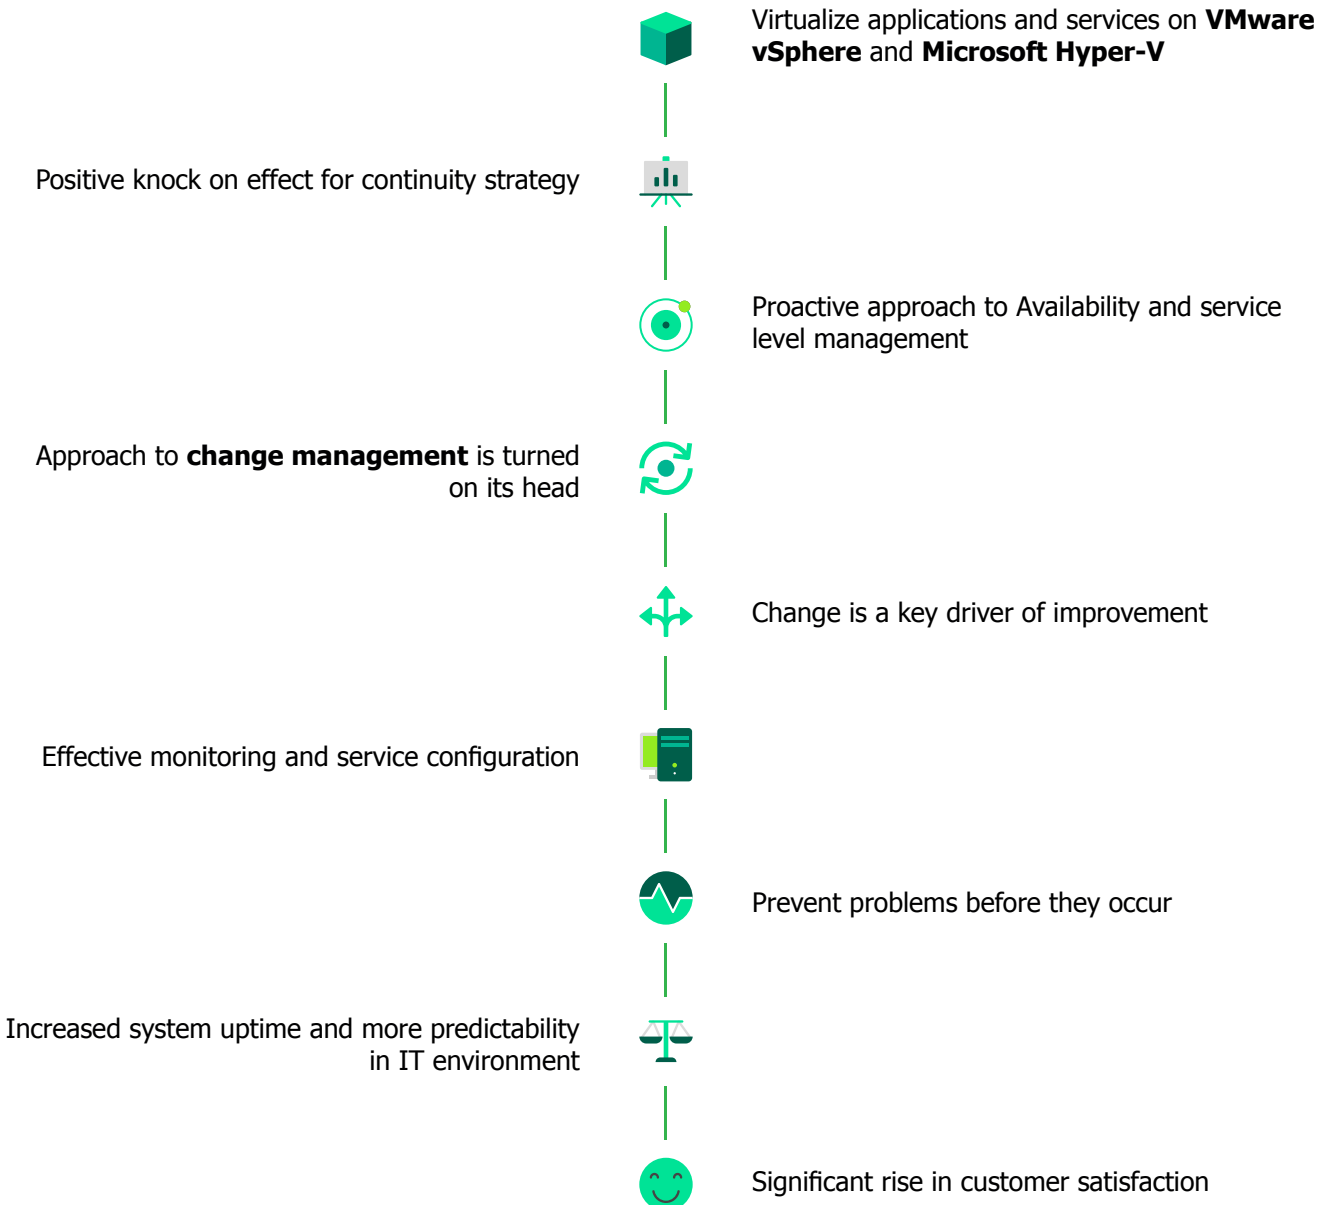
01

What You Will Learn

INTRODUCTION	**SETTING THE STAGE**	**IMPROVE CONTINUITY**	**MANAGING CHANGE**	**BETTER THAN A CURE**	**CONCLUSION**
What You Will Learn	The Fictious Company	Managing Availability and IT Service Level	Moving Beyond Reaction	Prevention and Continuous Improvement	Moving Beyond ITIL?

Aperature Labs' approach to ITIL

Virtualize applications and services on **VMware vSphere** and **Microsoft Hyper-V**

Positive knock on effect for continuity strategy

Proactive approach to Availability and service level management

Approach to **change management** is turned on its head

Change is a key driver of improvement

Effective monitoring and service configuration

Prevent problems before they occur

Increased system uptime and more predictability in IT environment

Significant rise in customer satisfaction

Aperature Labs, the Fictitious Company

Our company, Aperature Labs, is an American agricultural chemical manufacturer. Founded in Columbus, Ohio in 1988 by graduates of Ohio State University, its main focus is on producing pesticide products with minimal environmental impact. Aperature has enjoyed steady growth since it was incorporated. From its headquarters in Columbus, it has branched out to open laboratories and production facilities in Phoenix, AZ; Calgary, Canada; London, UK and, most recently, Perth, Australia. Employing around 150 people across all sites, and with a turnover in excess of $400 million, Aperature has a truly global reach, with a key focus in the APAC, EMEA and American markets.

Over the course of two decades, the technology available to the chemicals manufacturing sector has changed considerably. Rivals began making effective use of data mining and advanced analytics to improve process efficiency and predict demand more accurately. They started to use robotics and artificial intelligence as next level automation arrived. Crucially, they also used **new technology to improve service**, for example by allowing customers to book orders online through ecommerce platforms.

The IT Struggle of Aperature Labs

As with any dispersed global enterprise, development of a robust IT infrastructure has been crucial for keeping operations running smoothly.

Aperature's vital network assets include production line software and hardware, a supply chain ERP system, payment processing platforms, and integrated company-wide telecoms. Because of the need to support several remote sites, the infrastructure requirements are complex. The company relies on high Availability across its network for every aspect of its operations, from production and engineering to process management and ERPs.

Aperature has never been opposed to change, and has always shown a desire to be at the cutting edge of what it does. To manage its critical IT infrastructure, it adopted ITIL. But was not without its pain points. Like many companies, Aperature Labs focused on *reacting*, rather than *planning*, and its approach was affecting operations.

Key challenges:

• At a time when competitors were achieving major service improvements through digital transformation, Aperature was struggling with **downtime and late deliveries.**

• Known **errors and data loss** were causing instability.

• The IT department did not feel it had the tools needed to make meaningful **capacity plans.**

• Automated **monitoring alerts were unpredictable** – and off the record, IT would admit that the alerts mailbox was not checked as frequently as it should have been.

• Aperature did not have a full **Service Continuity**, nor a **Change Management process** in place, so software errors were not being proactively dealt with.

• Staff were using their own workarounds, with exactly the results you would expect: **a degree of chaos**, and a whole lot of re-work. The business was essentially too busy fire-fighting to make any positive progress.

Let's see what happened.

Managing Availability and IT Service Level

How Available is "Always-On" Availability?

Three nines? Four nines? 100 per cent? Is that even possible? How about: as available as your customers need you to be?

IT is fundamental to how most businesses operate. System Availability is a critical factor – even for normal, everyday functions. This is especially obvious in a manufacturing business like Aperature's.

Downtime in any part of its IT infrastructure could lead to it not being able to process orders or arrange deliveries, but also to not being able to produce stock. None of this was lost on Aperature. It was part of the reason it embraced ITIL in the first place. But like so many early adopters, it was missing a trick: what system Availability could mean to its business. It viewed the goal of Availability management as simply keeping the lights on – if a problem arose, it was fixed. If there was no problem, everything was fine.

Aperature had three problems to address before it could embrace **Availability as a key driver of the business:**

1. It did not have the right backup and recovery tools in place to get a grip of Availability.

2. Its monitoring systems were out of date and erratic, so it was never able to get the complete picture of what was causing system outages.

3. It was not able to anticipate issues before they arose, and work proactively to keep systems online, so there was continuity of service across the business.

Things started to change for Aperature when it took steps to integrate core systems across all five of its global sites. Let's look at how all that led to a step change in how they managed Availability, as well as IT Service Continuity.

The Challenge: Adopting New Systems

Relying on its distribution partners in the channel to get its products to market, it knew that one of its vital business functions (VBFs) was the speed and efficiency with which it could fulfil orders. The company therefore invested many years ago in ERP software to help improve the end-to-end integration of its production, inventory, sales and accounting processes.

As all good students of ITIL know, the service lifecycle does not end with the launch of a new system. This is where the people responsible for managing the infrastructure – in ITIL language, the IT service providers – must earn their corn.

So, with the ERP system up and running, Aperature's IT people kept talking to the business:

- Is the system efficient enough?

- Are orders being turned around quickly enough?

- Are we meeting service levels enough of the time?

The answers gave a mixed message. Maybe things could be better, and the IT team realized why.

With ERP platforms installed on data centers at each of the company's five offices, maintenance and service times were effectively being multiplied by five. Whenever there was downtime at one site, there was a knock-on effect through the whole chain of supply, from production through sales, hurting how efficiently orders could be fulfilled.

ERP Systems

Sales	Finance	e-Commerce	Warehouse Management	Purchase	HRMS	CRM	Inventory Management

03 Managing Availability and IT Service Level

The Benefits of Virtual Availability

Time to step up. The service owners took it upon themselves to find a way to improve Availability of the ERP system across all five sites. They started running all of the ERP software platforms in a virtual environment.

They also drew up an **Availability plan** based on **core business objectives** and demonstrated that **virtualization:**

- Would **reduce downtime** by providing valuable contingency.

- Would **improve disaster recovery protocols** by ensuring virtual images of faulty applications could be redeployed almost instantly, reducing the recovery time objectives.

- Would also provide a basis for dynamic capacity management to **optimize performance** at peak time, if lady luck was shining on the business.

This is a model of ITIL best practice: deploying IT in the service of clearly defined business needs, and taking steps to improve the service, with processes and their contributory factors well controlled at an organizational level.

And it all went beautifully. Until something went wrong.

Unpredictable Trouble

One day, a disk error caused a database in the system to be corrupted.

In the normal run of things, the centralization of the virtual ERP platforms worked in Aperature's favor, reducing maintenance downtime and allowing system management to be more efficient. But on this occasion, it conspired against the company; everything was brought crashing down. The only solution was to roll back the database to the last saved version from a day earlier. That meant a day's worth of data - sales, inventory, invoices - was lost, and had to be re-processed from paper records. This slowed down order fulfillment, and distribution partners were left without stock they required.

It was, in the classic sense of the word, a disaster. The worst system failure Aperature had seen in years.

When it comes to agreeing an IT service level, terminology matters. And, as we have already mentioned, so does the intention behind the service level – is it just about hitting the figures, or is everyone aware of the bigger picture?

So maybe it went a little something like this:

- The business felt that the amount of **downtime caused to the ERP systems** by carrying out maintenance five times over at five different sites was slowing down processes.

- The tech people took them at their word and came up with a solution which, in their language, drastically **reduced Mean Time Between Failures (MTBF)** - turning five lots of maintenance and service requirements into a single, virtualized system.

- No one realized what would happen if there was a critical failure to one of the now centralized databases. Outstanding Availability is not just achieved through increasing the mean time between failures, but also by **reducing the Mean Time To Repair (MTTR).**

The 24 hours it took to manually re-input the lost data showed there was still room for improvement.
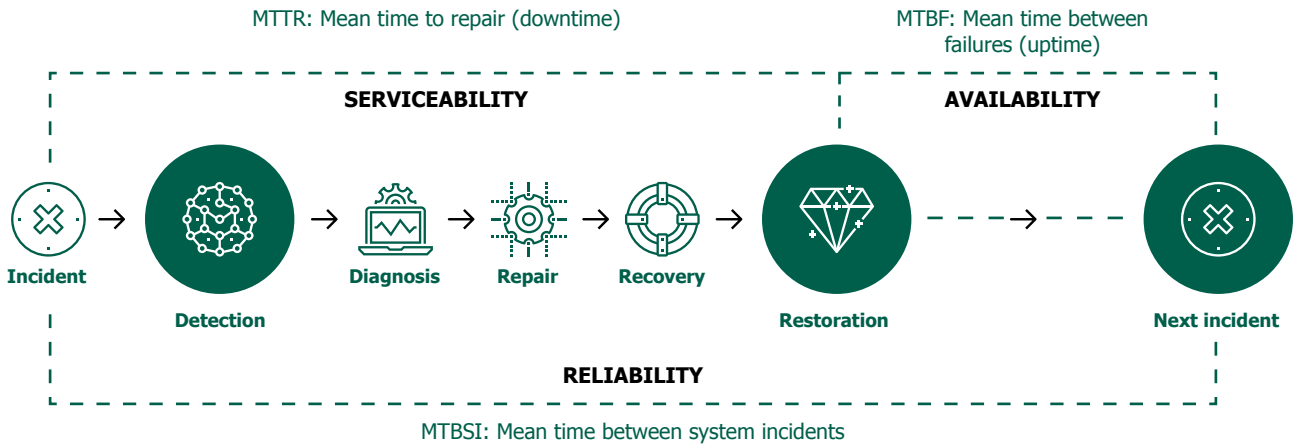
Managing Availability and IT Service Level

ITIL Incident Life-Cycle

MTTR: Mean time to repair (downtime)

MTBF: Mean time between failures (uptime)

SERVICEABILITY

AVAILABILITY

Incident

Detection

Diagnosis

Repair

Recovery

Restoration

Next incident

RELIABILITY

MTBSI: Mean time between system incidents

Lessons Learned

Fortunately, Aperature has the right kind of culture to be able to learn from its mistakes. This proved to be a transformative lesson in not simply following the letter of ITIL protocol on Availability. There's nothing quite like 24-hour delay to order fulfillment to remind you that service to the customer is the priority.

One of the big advantages of a using a virtualized environment across a dispersed IT infrastructure is how much easier it makes oversight and recovery.

Determined not to risk another significant disruption to its commitment to deliver outstanding service on order fulfilment, Aperature:

• Continued again on the improvement cycle, focusing on correctly documenting the procedures required to correctly failover its virtualized ERP.

• Rolled out virtualization across all its core IT functions.

• Installed a suite of software tools which combined backup, replication and restore functions with advanced reporting tools, promising recovery of all applications and data in under 15 minutes – a whole lot faster than 24 hours' manual work.

Service is Everything

Let's go back to the idea of viewing Availability in relation to your customers. We talk about **Availability targets as Service Level Agreements (SLAs)**, but is chasing a fixed figure really the right approach? At its logical conclusion, service is about putting the customer first, and doing whatever needs to be done to meet their demands. With competitors waiting in the wings to offer a better service, you cannot afford to be static.

Thinking about the dialogues which take place between the business people and the tech people in a company is fertile ground for exploring where missed opportunities in IT service delivery occur. We can imagine that, for the business side of an organization, delivering the best possible service to customers is always front of mind. The question is how that gets translated over into the practicalities of IT.

The example of Aperature's move to a virtualized ERP architecture demonstrates what can happen when an organization lacks the capabilities to see the big picture from an IT perspective. Yes, the goal was to cut out duplication in maintenance downtime, and therefore have higher Availability to meet the expectations of distributors and customers more efficiently. But what was actually delivered ended up not meeting the business objective, and it took a major system outage to get things right.

In ITIL terms, we can say that the SLA for the switch to a virtualized ERP, initially at least, did not meet the service requirements. Something major was overlooked. This demonstrated that Aperature was at this point still immature in its adoption of ITIL, as its Service Level Management was not satisfactory.

INTRODUCTION
What You Will Learn

SETTING THE STAGE
The Fictious Company

IMPROVE CONTINUITY
Managing Availability and
IT Service Level

MANAGING CHANGE
Moving Beyond Reaction

BETTER THAN A CURE
Prevention and
Continuous Improvement

CONCLUSION
Moving Beyond ITIL?

It was not always successful in turning business objectives into IT solutions.

It's easy to assume that technology alone delivers IT solutions. Not so.

Central to the idea of good **Service Level Management** is:

a) You must know what the service requirements are.

b) You must understand the capabilities of the technology you are using to meet those requirements.

By weighing these two factors in balance, you set parameters for how your IT system will work. These parameters form your SLAs.

When Aperature overlooked the MTTR in its newly virtualized ERP system, it did not fully understand the capabilities of the new technology. The failure to set robust parameters on what the IT could do led to a weak point. This is a very common issue. And without robust SLAs for their IT, organizations don't quite seem sure what to do with problems. Issues are left unresolved, creating a drag effect on the whole system and affecting continuity.

Even after Aperature virtualized its infrastructure and introduced **backup, recovery, restoration and reporting** tools;

*   It wasn't immediately able to deliver the **high levels of Availability and continuous IT service** it aspired to.

*   Its recovery and reporting tools promised **failover of under 15 minutes for all systems**, but the company was still a long way short of that.

*   It had the tools to give it effective oversight of the whole IT network, churning out **large amounts of performance data**. But it did not yet know how to use that data to turn it into meaningful action.

The Exchange Server Outage

Aperature had been coping with a relatively longstanding issue with its Microsoft Exchange server. Employees had been reporting problems sending emails for some time. Incidents built up; nothing was done.

Management decided it wanted to get a handle on how incident reporting was being addressed. It asked for a full report on where all the pain points in the IT infrastructure were.

It found numerous problems:

*   The team had found it difficult to cope with the large number of automated alerts created by the new system.

*   It was unable to respond to every single alert as an incident, and recognized that many were trivial in isolation.

*   Data was stored in separate silos server by server, so creating an accurate comparison across the whole system would be time consuming.

But the issue here was bigger than the Exchange server. This was a classic example of there being no meaningful SLAs in place. No-one had considered the alerting process. No-one had thought about keeping this data in a unified format. Without clear thinking, the potential to drive genuine service improvements was missed.

Aperature's management agreed to upgrade its reporting software, choosing a product which creates a visual heat map on a GUI which instantly show where alerts are being raised. These heat maps quickly showed issues that the remote sites were causing bottlenecks. Without the wide scale visualization of how the network was functioning, this would have been very difficult to spot.

Four Takeaways from Chapter 3

Installing and managing a **backup & recovery solution** on a daily basis reduces risks for your IT infrastructure and **ensures business continuity.**

Full system oversight is critical to **improving Availability** and continuity in IT systems, but must be accompanied by robust SLAs setting out how the data is used.

IT system **SLAs** should reflect **real business performance objectives**.

Clear communication and understanding between **business units and IT units** is critical for delivering against business objectives.

Change management in ITIL is all about mitigating risk, not leaving you exposed to more incidents

So far, we have followed Aperature as it has taken some significant steps forward in its IT management. We have seen:

- The introduction of centralized ERP system in a virtual environment to make control and maintenance more efficient.

- Rollout of a virtualized architecture across its whole network.

- Dedicated cross-system backup, recovery and replication tools to improve Availability and IT service continuity.

All proactive and positive moves. But we have also seen a flipside to this.

Change leads to more incidents, some of them more critical than the problems which had gone before. Per the Forrester report, changes in IT systems are the top cause of incidents in business networks, cited by 78% of respondents. Expansion in scope of services was a close second, with 72%.

Just as introducing change without robust management and planning leaves you exposed to the risk of more incidents, so does avoiding the risk altogether and not doing anything at all.

Change management is supposed to be the antidote to that. Done properly, change should be about reducing incidents, stopping the cycle of having to go back to the same problems over and over again, and preventing new ones occurring. And **change management can and should be the one key way of moving beyond reacting to incidents, and into the realms of proactive prevention**.

We shouldn't downplay the importance of fixing problems. But if the same incidents recur, a change is required.
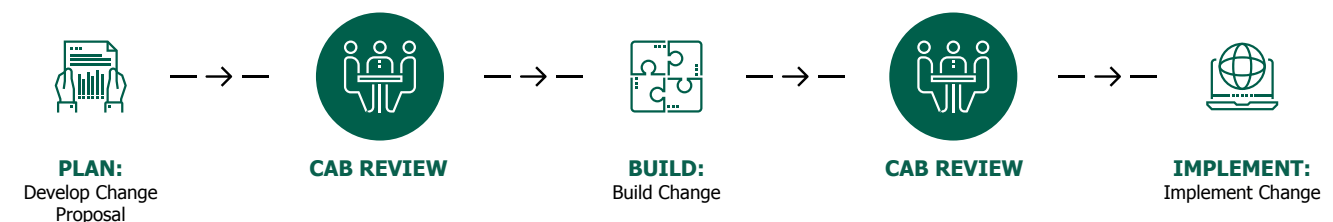
Improving change controls

Aperature had a known problem with production line software crashing. The IT service team was quickly able to work out that there was a problem with the way their innovative scales used to weigh out materials were configured. When items were removed from the scales having been weighed, the reset to zero would sometimes cause a software error and the whole application would crash, putting valuable design and engineering data at risk.

The IT team developed a workaround: reboot the server. But Aperature was **counting the downtime** it caused in hours rather than minutes.

> Introducing change without robust management and planning leaves you exposed to the risk of more incidents.

Change Management in ITIL

PLAN:
Develop Change Proposal

CAB REVIEW

BUILD:
Build Change

CAB REVIEW

IMPLEMENT:
Implement Change

Change controls started to improve once Aperature decided to reduce risk of data loss through, backup and recovery software and add monitoring and real-time alerts tools for their backups and virtual infrastructures, as well as use testing environments for some of their upgrades:

• The backup and recovery platform **reduced risk of data loss** through image-based backup, advanced replication and backup from storage snapshots.

• Aperature also started to put in place a **streamlined DR** with the built-in tools, and be prepared for any inevitable disruption caused either by human error or the new technologies they started implementing.

• **Real-time monitoring** and alerts were put in place to notify with alarms of backup and VM performance issues, to avoid downtime and meets SLAs.

• By working in a **replicated virtual lab**, the reconfigured VMs could be tested and verified before being launched in the production environment.

Change management was taken one step further by putting in place a **Change Advisory Board (CAB)** for its IT systems management, to assist and prioritize changes and in the same time keep production seamless and uninterrupted. Also, **drawing on the Seven Rs of Change Management**, the IT team built a stronger case. Incidents were correctly recognized into Problems.

This was a watershed moment. For the first time, Aperature had gained the confidence to embrace change. Not only could they identify weak points and make improvements before weaknesses escalated into endless streams of incidents, but also made sure they were prepared for disruption. As an ITIL practitioner, Aperature was now on its way to **embracing continuous improvement.**

Four Takeaways from Chapter 4

Many businesses experience spikes in downtime and incidents following changes to their IT system.

Effective change management should anticipate and plan for such incidents so they can be avoided.

Good backup and replication tools support effective change management by helping to avoid mitigate against incidents arising from change.

Being prepared for the disruption change can bring by adopting Availability solutions, DR strategy, as well as BC planning will, for certain, reduce risk and improve change controls along the process.

The Seven R's of Change Management

Raised
Who raised the change?

Reason
What is the reason for the change?

Return
What return will the change deliver?
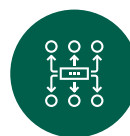
Risks
What risks are there if we do/don't carry out the change?

Resources
What resources will be required to perform this change?

Responsible
Who is responsible for this change being performed?

Relationships
What relationships are there between this & other changes

05 Prevention and Continuous Improvement

Recap: Why Are We Here?

Aperature's fictional journey through ITIL implementation shows one key truth. When it comes to IT management, prevention of problems is always better than a cure.

When we looked at Aperature's experience with Availability management, we argued that high levels of continuity depended on thorough system monitoring and backup and recovery capacity, but also a clear strategic vision for how performance data was to be used, prioritized and acted upon. More explicitly in Chapter 4, we made the case that fully embracing change management critical in Aperature's ITIL journey because it opened the door to a preventative approach.

Let's recap on the role of **IT performance as a critical function of business performance**. It's key to understanding why the concept of continuous improvement should enter the ITIL landscape.

When we talk about IT systems, or approaches to IT management like ITIL, it is easy to get lost in technicalities. But we need to remember that we do not have IT systems in our businesses for their own sake. The business is the ultimate benefactor.

Achieving high system Availability is about meeting expectations. Proactive management and real-time oversight achieves the same thing. In this chapter, we will delve into examples of what a preventative approach to ITIL looks like in practice, explaining how Aperature became truly proactive in **optimizing its IT systems.**

The Goldilocks Principle

How much capacity do you need in an IT system?

Not too much. Not too little. Just the right amount.

Capacity management is all about **striking a balance between performance and efficiency**, because if you don't leave enough headroom, spikes in usage cause slow speeds or downtime. You wind up fighting fires. If there is too much headroom, you are writing blank checks.

Being responsive is important. We all know capacity demands on IT infrastructure fluctuate by season, or in response to marketing campaigns, or because more people logged on that day. If you want to get capacity management 'just right', you must be able to ride the waves.

Aperature wanted to find its own Goldilocks point for two reasons:

1. To maintain its commitment to delivering outstanding levels of service, regardless of demand.

2. To control TCO of its IT infrastructure.

Virtualization was a key enabler in achieving this.

The company developed a capacity plan based on its newly centralized architecture built around the software-defined database at its Columbus, OH headquarters. The VMware vSphere platform Aperature ran offered flexibility. Dynamic resource scheduling meant server resources could be reallocated whenever additional compute capacity was required. With proper notifications, the

Capacity Management
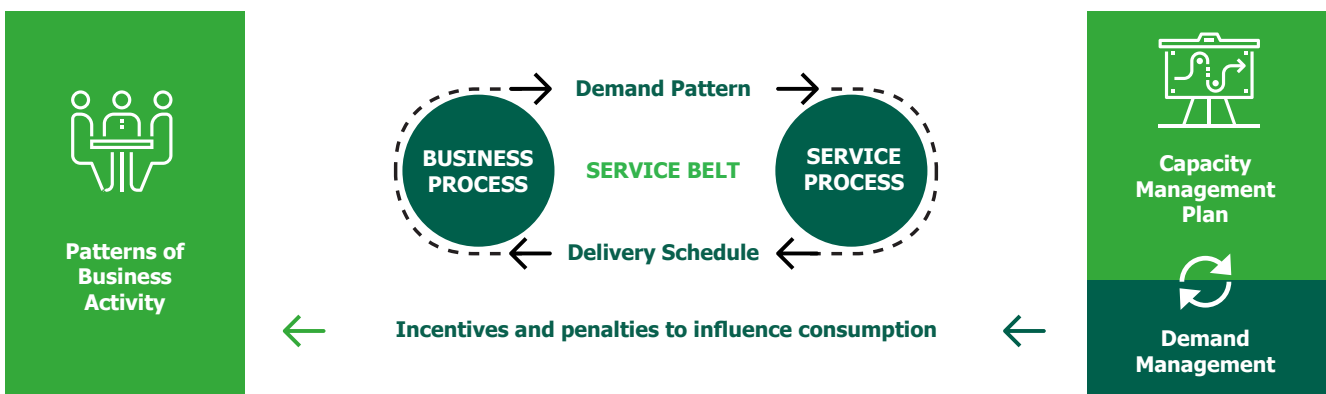
Prevention and Continuous Improvement

capacity plan made provision for new servers to be brought online before issues occurred.

But at first, Aperature still found itself firefighting capacity issues. It had not yet fully grasped all the implications of its system changes, and was running into new incidents it had failed to anticipate.

The following summer, Aperature found its products more in demand than ever. Core services – like ERP – stood solid. Availability was not an issue. But capacity was pushed to the limit. VoIP calls began to suffer quality issues as packets met competing network traffic. Remote lab sites saw slow network speeds as connections to these remote servers were stretched.

Aperature had based provisioning on historical data:

- It was not ready for an unprecedented spike in traffic.

- Its service capacity management was flawed.

- It got its predictions wrong.

- It was forced to make unplanned emergency changes on the fly.

- It was still reactive, and leaving the system exposed to more problems because these emergency changes were largely untested.

- It was completely unprepared for a spike.

Dynamic demands require dynamic oversight. Aperature quit planning capacity based on what had happened a year ago. Instead, it configured its monitoring tools to:

- Provide day to day visibility across all IT environment from physical to VMs detailing network traffic, usage, system outputs and backup jobs.

- Provide real-time alerting of backup and infrastructure performance.

- Obtain predictive power - patterns merging over days, or hours, allow you to forecast system requirements and adjust accordingly.

With the right component management, supported by advanced reporting tools, Aperature was therefore able to start acting proactively to prevent capacity issues.

Putting All the Pieces Together for Continuous Improvement

If you want to take IT management beyond reactive firefighting and towards optimal performance, proactive approaches are essential across the board. Only then can you harness IT to deliver on your business objectives of always striving to deliver better service.

In ITIL terms, **continuous improvement is about always making small adjustments to ensure the maximum possible Availability and continuity.** No issue is inevitable. Service levels are not fixed. And change should be embraced as part of the culture which drives the organization's ambitions.

Let's look at one final area of Aperature's ITIL journey, which serves as a shortcut for how it turned reaction into positive action using good analytics tools.

Configuration Management: The Domino Effect

Our journey reveals a fundamental truth: everything is interconnected. If one piece of the jigsaw does not fit, it has a knock-on effect overall.

We saw how a problem with scales hardware was causing production software to crash, and how the failure of a single database in the newly virtualized ERP environment brought down the whole system. In an efficient, well optimized system, everything has to be carefully configured to be able to run at its best in sync with everything else.

In ITIL, this is what Service Asset and **Configuration Management** teaches us:

- Every platform, every application – what is known as a Configuration Item, or CI – works in relation to everything else. In other words, you need to know exactly what will happen if one CI goes down, what impact it will have on the next item it is joined to in the network, what the broader impact on system Availability and performance will be.

05 Prevention and Continuous Improvement

- To be successful, you need to have good insight into how the process is operating. And that is not just a snapshot here, and a snapshot there. You need a continuous, ongoing flow of data-led intelligence about your system configuration in real time. Anything less, and you will always be left chasing your tail.

- And you need good analytics to be able to continuously improve.

Change Management
Assesses Impact

Capacity Management
Assesses Impact on Business & IT Performance

Configuration Management
Identifies Areas Impacted

We have already seen, for example, how Aperature's decision to base capacity planning on historical data led to a shortage when business suddenly spiked and the systems were inundated with orders, requests, calls and processing actions. We also saw how decision making over how to handle automated incident alerts led to an Exchange server issue being left unresolved.

Good configuration management also has a human element. Items in a network are not the only things which have an impact on each other. Human inputs are critical, too.

All of the company's labs, in Phoenix, Calgary, Perth and London, as well as at company headquarters at Columbus, OH, all use the same virtual servers to log results by project. If they are running tests for one product development project, they share results on one server. If they are running production analysis for another, they record data on another. One day, in preparation for a new project, a member of the service desk team logged on. He noticed a few unused virtual machines lying around, and in the interest of being neat and tidy, decided to delete them.
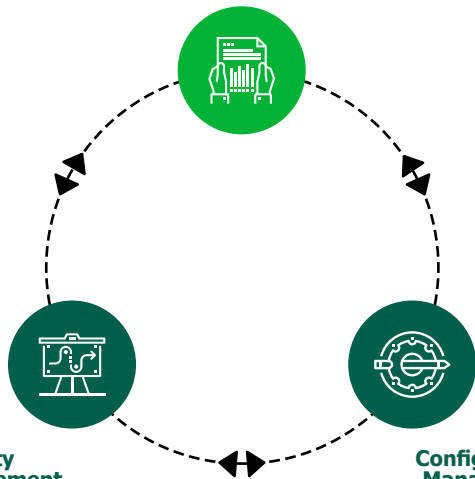
Bad move. One of the deleted VMs was still being used by a legacy system. Equipment stopped working.

Aperature's monitoring system created a ticket. As the reporting software had full auditing capabilities, it identified where the change (or in this case, deletion) in the virtual environment took place, exactly what had been changed, and what its impact on the rest of the system would be (the legacy platform which no longer worked).

Knowing exactly what had been deleted, the service team could simply go to the back-up and replication environment, find the copy of the VM they needed, and restore it. The incident was resolved before it even became a problem.

> With advanced analytics tools, reporting gives you predictive power. Patterns merging over days – or hours – allow you to forecast system requirements and adjust accordingly.

05 Prevention and Continuous Improvement

Lessons Learned

Why would you delete VMs without running a check to see what they were used for? This is only likely if the team doesn't understand how the virtual layer in which he was working related to the application layer.

In this sense, many examples of human error in IT management can be attributed to a failure of configuration management. Or to flip it around, **good configuration management is a way to mitigate against the impact of human error on an IT system.**

Here's the key point. Configuration management does not work as a reactive response. ITIL orthodoxy teaches that IT systems are intrinsically dynamic. Configuration management must therefore be approached as a process.

Given the huge amount of complexity which is now built into business IT systems, such as the virtualized infrastructures operated by globally dispersed businesses, this is very difficult to achieve manually. But with effective, continuous, automated whole system monitoring and analysis in place, it can be simple to achieve.

Had Aperature still been in the situation where it took static snapshots of system configuration and performance, it might not have spotted the issue until days or weeks later, and then would have been scrambling to react as something simply did not work.

Four Takeaways from Chapter 5

Preventing problems rather than reacting to them requires an approach built around continuous improvement.

In-depth system analytics are an important tool for informing the predictions on which you base proactive preventative measures.

Real-time reporting also allows issues to be identified and resolved before they have an impact on performance, maintaining continuity.

Automation in reporting and failover support helps to mitigate against the effects of human error.

> Good configuration management is a way to mitigate against the impact of human error on an IT system.

Moving beyond ITIL?

Shifting IT Operations from Reactive to Proactive

We've followed Aperature on a journey to find its feet with ITIL. Like so many companies from midmarket up to enterprise level, Aperature engaged with ITIL because it recognized its IT systems as an essential asset. IT is fundamental to production, supply chain management, distribution, customer service, sales, communications, and coordination of complex tasks across sites.

Aperature saw that if it wanted a lean, responsive, productive business, it needed a lean, responsive, productive IT infrastructure too. ITIL is recognized as one of the leading strategic approaches to organizing and managing IT systems in a way which delivers genuine business benefits.

But Aperature's journey underlines that ITIL is just the start.

When we first met Aperature, it was already well on the road with ITIL and looked to be doing the right things. There was good communication, there were SLAs, there was planning. It switched from on-premises data center deployment across five sites to a virtual environment. The new system was designed and built to improve efficiency, centralize control, and allow for better capacity planning. Moreover, the tech people were listening to what the business people wanted, and working to deliver for their needs.

But something was still missing. Aperature was reacting. The service desk was in permanent triage mode. It remained risk averse when it came to considering change to improve performance on some vital business functions.

Adopting the right tools was key to Aperature being able to move from a reactive to a proactive stance in its ITIL implementation. Here's why:

- No network is static. IT systems are - by definition - a series of interdependent, complex processes. N inputs cause new effects further down the chain, and things go wrong.

- Your approach needs to be based on continuous process-based monitoring, flexibility and anticipation. When you don't get it quite right, you are left chasing your tail.

Looking back, we can identify 3 key ways that Aperature took ownership of its IT problems and matured its approach to ITIL using tools:

1. It introduced a robust backup, recovery and replication platform, which transformed Mean Time to Repair in its virtual environments from hours into minutes, also mitigating risk of downtime and the level of stress in its IT professionals. Also, with built-in DR capabilities, they managed to put in place a streamlined DR strategy.

2. It installed a monitoring, reporting and capacity planning suite, which automated its incident alerts and ticketing and gave it the ability to forecast resource usage and utilization trends.

3. With increasingly well-defined processes and procedures, and the use of statistical data to control them, Aperature began to respond to incidents in close to real time, putting out fires before they had a chance to have an impact, and using change in their own benefit, to stop anything happening in the first place.

Key Benefits of ITIL

Support business outcomes

Enable business change

Manage risk in line with business needs

Continually improve

Show value for money

Optimise customer experience

06

Moving Beyond ITIL

From ITIL, Where Next?

Is it enough to follow ITIL as a fixed set of principles for IT management?

If you really want to get the best out of an IT system, does there need to be a more fundamental shift in how we view IT system management as a discipline?

Perhaps the real issue is one of interpretation.

Like Aperature, many businesses say they are following ITIL when they are reacting to incidents. But minimizing incidents is just as important. Aperature's watershed moment was realizing that it could make changes whilst also minimizing risk.

Supported by backup, recovery, replication and monitoring tools in its physical, as well as virtual environment, it could focus in on the root of changes, build and test new systems without having to take down old ones, and create failovers through replication should anything go wrong.

Change no longer seemed such a big deal. Change could happen as much as it needed to.

So instead of workarounds and hacks, Aperature was now finding root causes. This is flexible, predictive, proactive system optimization in action.

When people talk about IT system performance, everything is described in segments. It is built into the language of ITIL - the seven pillars, systems built up of individual service assets, service delivered through lots of different SLAs. But the big picture is the most important. Businesses are not judged by what all the individual parts do in isolation. They are judged on what the sum output of those parts can be.

When Aperature's management decided it needed whole system insight on IT performance, it was looking at its digital infrastructure in a completely new light. No longer was it a series of assets and incidents. It was an interconnected whole, supported by the right tools and technologies. Robust reporting and analytics, supported by GUI visualization, breathed life into that understanding. That is the final condition for optimizing a dynamic system - understanding how all the parts play into the overall performance. And this is where ITIL adoption comes into its own.

> ITIL is recognized as one of the leading strategic approaches to organizing and managing IT systems in a way which delivers genuine business benefits.

07 Glossary

Availability management:
Planning availability of all IT systems across the business

Change advisory board (CAB):
A group of people who meet to assess changes and prioritize them

Capacity management:
A process of ensuring that systems can cope with peak demand without overspending

Capacity plan:
An analysis of different capacity requirements along with the cost of achieving the required capacity

Change management:
A process of moving a system change from planning (Service Design) through to live (Service Operation)

Configuration items (CIs):
The hardware and software in the business, plus documentation or other resources

Configuration management database (CMDB):
A database containing all the configuration items in the business

Enterprise resource planning (ERP) system:
A system where all major business functions are managed

Exchange server:
A server that handles Microsoft Exchange data, such as email and calendars

GUI:
Graphical user interface

ITIL framework:
A set of best practice guidelines for using ITIL in your business

Mean time between failures (MTBF):
The mean time that elapses between each IT failure

Mean time to repair (MTTR):
The mean time that elapses between an IT failure and the repair

Microsoft Hyper-V:
A hypervisor for Windows that allows administrators to create multiple virtual servers

Request for Change (RFC):
A document provided to the Change Advisory Board containing details of a potential change

Service level agreements (SLAs):
A contract that sets out the level of service that a customer should expect to receive

Seven Rs of Change Management:
A list of 7 criteria that are used to assess change: who Raised it, the Reason, the Return, the Risks, the Resources, who's Responsible, and any relevant Relationships.

System availability:
The amount of time that a system is online and functional

Vital business functions (VBFs):
Essential functions within a business that are supported by IT

Virtual environment:
A self-contained system deployed on a host. Several virtual environments can co-exist on the same computer.

VMware vSphere:
VMWare's virtualization platform

Voice Over IP (VoIP):
A telephone system that uses the internet to place calls, rather than a traditional PBX

08 Bibliography

Checklist Capacity Plan ↗

Downtime, Outages, and Failures: Understanding their True Costs ↗

Incident Priority in ITIL Explained with Examples ↗

ITIL Foundation Handbook (Stationery Office, Claire Agutter) ↗

The Seven Rs of ITIL ↗

The 5 Stages of ITIL in the Modern Enterprise ↗

Forrester Custom Technology Adoption Profile, "From Reactive To Proactive: The Age Of The Customer Demands Moving Beyond ITIL", January 2017

What to Expect When You're Expecting ERP System Downtime ↗