



WHITE PAPER: Cloud Computing

Key Service Management Processes for Cloud Infrastructures

Denis Nothern
Vice President, Application Product Design

Table of Contents

- 2 Introduction
- 2 Application Infrastructure Evolution
- 4 Key Service Processes
- 9 Summary
- 9 About Savvis

Introduction

Cloud computing has significantly impacted how companies deliver and support applications. On the positive side, cloud computing has the potential to both increase end-user productivity and reduce infrastructure costs. But a less-desirable byproduct of the shift to cloud computing is that IT support teams must manage increasingly complex applications and infrastructure. This paper will systematically consider how the advent of cloud infrastructures has affected seven key IT service management processes using a hypothetical case study drawn from Savvis' customer experiences. We will follow the evolution of application infrastructure at "Folio Management Corporation" and then describe the company's deployment of these seven processes as FMC moves its applications into the cloud. We will also explore how these processes can be improved by automating deployment of cloud infrastructure.

A number of factors have caused application platforms to grow in complexity, including the move from centralized computing to highly distributed computing; the transformation of physical infrastructure to virtual infrastructure; and the increase in client platform alternatives. Traditionally, growing complexity leads to greater risk and higher costs due to the greater number of components that companies must control. Because of this, IT infrastructure managers must first determine which IT processes need to be implemented in a cloud computing environment.

In the mid-1980s, basic IT processes were organized into systems management processes that depended on standardized steps and procedures to accomplish specific tasks. Over time they evolved to include process performance standards that focused on the quality and reliability of the work being done, and process management standards that focused on measuring and regularly improving the targeted business processes. In recent years, IT best practices and processes have been well documented in the IT Infrastructure Library (ITIL). Version 3 is the most recent version of this widely accepted IT resource.

Application Infrastructure Evolution

Launched in the 1980s, Folio Management Corporation (FMC) develops and sells portfolio management products for small and medium-sized (SMB) financial services companies. The founders were business and application architects who had worked at larger companies where mainframe computer applications were used to manage thousands of investment portfolios. However, the age of mainframes seemed to be ending. FMC's founders believed that demand for centralized mainframe applications would begin to rapidly decline due to the increasing costs of supporting such a platform as well as the recent introduction of the client-server architecture. They designed a new application that leveraged the increasing intelligence of end-user workstations. By distributing the work a centralized mainframe-based application typically performed between a minicomputer and these client workstations, they were able to deliver a much higher-performing and lower-cost application across a broad range of scenarios (Columns 2 through 6 in Figure 1 of next page). From an architectural perspective, the application was similar to the mainframe version in that most of the application still ran on a central processor. However, the application provided end users with significantly greater ease of use and flexibility.

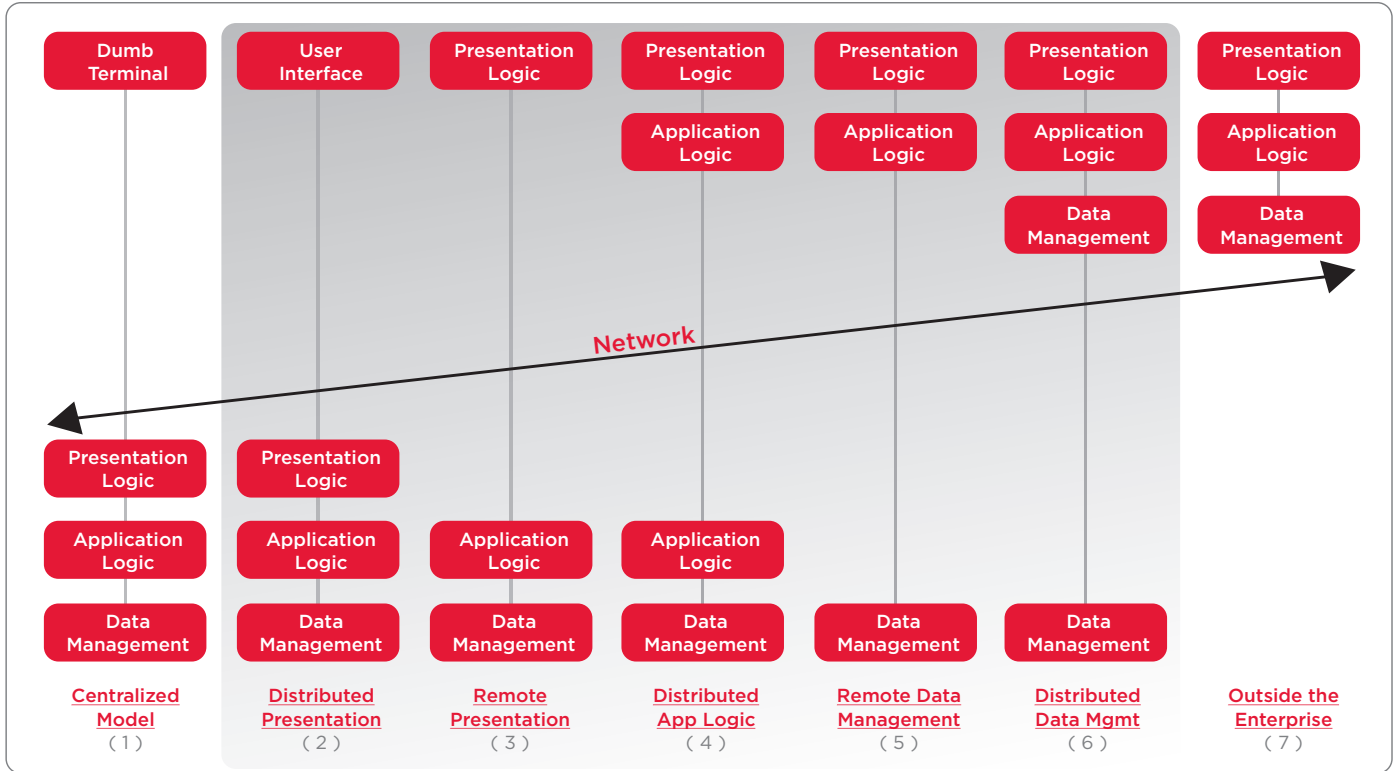


Figure 1

Then, in the mid-1990s, the Internet revolution hit. FMC faced the first major redesign of its application. In response to customers who wanted their portfolio managers to access the application from field-based laptops, FMC reengineered its application to run on a standard server infrastructure that was distributed across a three-tiered architecture. The new version also introduced a distributed database design that allowed a portfolio manager to have a thick client installation on a laptop and to synchronize changes with the central database every few days. This transition was not without its challenges, however. FMC realized it no longer had tight control of application infrastructure, which forced its development teams to spend extra time and effort designing and testing its application for use on various servers and client systems.

In the early 2000s, the architectural landscape changed yet again. By 2001, progressive FMC customers no longer wanted to own and support the application's infrastructure, and pushed FMC into offering the application infrastructure as an Application Service Provider (ASP). This proved to be more cost effective for FMC customers, since they no longer had to spend capital on back-end equipment — only on client machines. Over time, however, this approach proved to be even more complex for FMC support teams since they now needed to test the application — and any changes they made to it — across multiple combinations of hardware and software.

WHITE PAPER: Cloud Computing

FMC also realized it now faced more-frequent infrastructure architectural changes. When the software-as-a-service (SaaS) market started to develop in mid-decade, FMC kept pace, introducing a SaaS offering in 2005 and consolidating multiple customers onto fewer multi-tenant platforms.

It was at that point that FMC became interested in cloud computing. It began searching for a managed services provider that offered a cloud option. This would allow FMC to replace its fixed-cost application infrastructure with a variable-cost model that it could better align with its revenue stream. In 2009, FMC found a cloud computing solution with a service level agreement (SLA) that delivered the reliability and performance its customers expected.

The evolution of FMC's application infrastructure is similar to that experienced by many Savvis customers. To accomplish this transition from a relatively simple and closed environment to a flexible infrastructure capable of supporting clients' needs, companies like FMC depend on a few well-executed ITIL Service Design, Transition, and Operation processes. The next section looks at these processes; how to effectively apply them for applications in the cloud; and what additional benefits companies can gain from this transition.

Key Service Processes

By the time FMC made its move into the cloud, it had already recognized that application infrastructures were changing quickly and that it had little control over the infrastructures running its applications. Over its 25-year history, the FMC IT support groups had grown increasingly dependent on service management processes to keep pace with increased infrastructure complexity and decreased infrastructure control.

Initially, FMC relied upon Change Management and Problem Management to reduce risk and costs. This meant that FMC's service management was largely reactive, as it only scheduled software changes when needed and fixed problems as they occurred. After introducing a formal ITIL framework in 2000, FMC recognized it needed to transcend simple Change Management and Problem Management to address the growing complexity of its infrastructure. In response to this complexity, FMC focused on Release Management to control software releases. After moving to the Internet, it implemented Security Management, and after introducing the ASP application model, formalized Capacity Management, Service Level Management, and Service Continuity Management.

Taken altogether, these seven systems management processes have provided FMC with the framework it needs to successfully transition to cloud computing. The rest of this white paper defines each of these processes and discusses how they can be leveraged to move applications from either legacy mainframe or client-server architectures into the cloud.

Change Management

First and foremost, Change Management defines what type of change is required to reach a certain goal. It also involves executing the right change workflow for the type of change in question, and remediating the change if it doesn't meet its objective.

WHITE PAPER: Cloud Computing

FMC understood the importance of controlling changes when the company first opened its doors in the mid-1980s. At that time its approach was fairly simplistic in that it only tracked planned changes and left many of the details of managing change and recovering from failed changes to the development groups that were actually making the changes. As FMC transitioned its application to the three-tiered Internet infrastructure, multiple application and infrastructure IT groups needed to apply changes to the application in a production environment and coordinating changes across teams became critical. Some groups also submitted a disproportionately high number of emergency changes since these could be implemented quickly without waiting for the weekend “change window” that standard changes needed to be implemented within. FMC created a separate quality assurance (QA) environment that paralleled the production environment as a way to test changes and increase the chance that the changes would be reliably implemented.

When FMC virtualized its Web and Application tiers for its ASP model, it achieved the expected benefits of reducing provisioning time as well as overall infrastructure costs. However, it also faced increased complexity since it now needed to support the additional virtualization operating system software.

By outsourcing its infrastructure to a managed services provider that offered cloud capabilities, FMC was able to implement change management processes much more cost-effectively and efficiently. Since the provider used automated systems and network platform provisioning, FMC could depend on the provider to build consistently standard and identical platforms.. And because its outsourcing vendor also used a configuration database that provided a clear view into infrastructure dependencies, FMC could refocus its developers' attention on application changes and reallocate infrastructure support resources to other projects.

Problem Management

The objective of Problem Management is to minimize the adverse impact of errors and prevent the recurrence of incidents due to those errors. Problem Management encompasses two key processes: addressing one or more incidents as they occur (also known as Incident Management); and eliminating recurring errors for which root causes have already been identified. FMC implemented Problem Management first to track the number of bugs in each release of its own software products. It was also able to use the root cause analysis of incidents to identify problem patterns.

Fixing the sources of recurring problems allowed the company to improve its software release process by incorporating tests during QA. However, as FMC redesigned its application to be delivered via the Internet, it found the number and types of problems increased due to the wider variety of application infrastructures and end-user workstations that ran its software. To address these problems, FMC took three steps. First, it standardized the types of infrastructure components its software could be deployed on. Second, it increased its control over application and infrastructure changes by bundling all changes into a single packaged release. Finally, it specified which types of laptop and desktops could be used as client machines.

WHITE PAPER: Cloud Computing

Initially, these actions significantly reduced problems FMC customers were having with the application. By standardizing server builds, the FMC development teams could depend on a well understood production environment. This in turn enabled them to package multiple application changes into fewer, larger releases that could be more thoroughly tested in their QA environment. However, the growing variety of available end-user workstations prevented FMC from limiting which ones customers could use. In realizing this, they changed their application's interaction with client hardware to only use standard programming interfaces.

When FMC moved to a cloud infrastructure, it was able to leverage the benefits of configuration automation by taking a just-in-time approach to QA. By creating and operating system integration and production testing QA environments only as needed, it was able to dramatically reduce costs. In addition, the provider's use of provisioning automation reduced the number and variety of problems caused by manual system and network provisioning. Again, because of automation, FMC was able to reallocate trouble-shooting resources to other IT projects.

Release Management

Release Management governs both hardware and software additions and changes to the production environment by ensuring clear deployment plans are developed; that release packages can be built, installed, tested, and deployed; and that a knowledge transfer happens with users and operations to optimize use and support of the service.

When FMC first created new product releases, application development managers only considered bundling all related changes for major software releases. When application problems increased during its Internet deployment product redesign, product marketing demanded increased quality controls to address customer complaints. At that point, FMC management introduced a formal release management process to govern all hardware and software changes, and developers no longer viewed changes in isolation since release management forced development teams to consider what might be impacted with each change.

When the FMC product moved to a cloud infrastructure, this well-structured release management process continued to serve the company well. FMC was able to rapidly create standard testing environments matched to the production infrastructure when needed, and release the resources when testing was completed. This cloud feature allowed FMC to test each release for completeness and stability before promoting the bundle into its production environment.

Security Management

Unlike other service management process, Security Management is performed according to the overall corporate governance framework that guides an organization's assessment and management of risk. From the onset, FMC management recognized the need for their product to be secure. They defined a security framework designed to protect users who relied on the information in the application and networks delivering that information, with the goal of providing a safe way of doing business with other organizations. This approach was specifically designed to prevent security problems affecting information confidentiality, integrity, and availability.

WHITE PAPER: Cloud Computing

At first, FMC required an account and password to access its application and defined roles for each type of user. It also supported several customer single sign-on solutions to integrate its application security with what was available at a customer's site. This worked well for the minicomputer and ASP versions of its product. However, when offering a SaaS solution, FMC recognized it needed to augment security with a perimeter defense to limit general access to its application. It also added security between each layer in the three-tiered architecture to limit unauthorized access between layers.

When FMC moved its SaaS offering to a cloud provider, it was well positioned to protect its application. Recognizing it no longer controlled the underlying application infrastructure, it made sure its managed services provider would take responsibility for patching the systems and supporting infrastructure. Since FMC wasn't planning on keeping its systems integration and QA sites continuously active, it looked for a cloud provider that understood the challenges of keeping the virtualized application infrastructure patched even for inactive systems.

Capacity Management

Capacity Management ensures that demand for and supply of capacity are balanced. After all, if supply and demand are out of balance, this directly impacts cost of service delivery for application infrastructures. To perform effective Capacity Management, organizations need to monitor application infrastructures, end-user experience, and infrastructure utilization over time to gauge that sufficient capacity exists to meet the requirements of SLAs.

When FMC reengineered its application for a three-tiered infrastructure, customers were finally able to add smaller increments of capacity to address user demand. However, this was difficult to deliver with the ASP model, since each customer required its own, isolated infrastructure and data. FMC addressed this by creating a multi-tenant solution, in which it isolated customer data in the database and virtualized the Web and Application tiers of its solution so multiple customers could run on fewer servers.

When FMC began offering a SaaS model, it added end-user experience monitoring to track transaction response time. This monitoring solution automatically determined normal end-user response time and sent alerts to warn of potential service level failures, improving FMC's ability to find the root cause for an incident by pinpointing which part of the application was causing the fault. FMC also used monitoring for capacity planning and to determine which parts of the application would be impacted as it made database schema changes. When FMC moved its application to the cloud, it was able to further reduce costs. Every time a new customer signed a contract, FMC was able to quickly activate another instance of the application to provide needed capacity. As customers cut back or stopped using the application, FMC was able to immediately shutdown the virtual infrastructure and stop paying for unused capacity.

Service Continuity Management

Service Continuity Management combines business continuity requirements to support the users of an application and IT service continuity requirements to address the hardware and software that supports the application service. As with Security Management, Service Continuity Management is typically part of a broader company business continuity effort with senior management participation.

WHITE PAPER: Cloud Computing

When FMC first offered an ASP solution, it realized it needed a recovery solution in the event of a prolonged service outage. Although it periodically tested its QA infrastructure as a replacement production environment should anything go wrong, this approach proved to be very disruptive for the change management and software release teams since they couldn't test new software when the infrastructure was being used as a replacement for production. It was also expensive to buy sufficient quantities of QA hardware as FMC's ASP customer base grew. When FMC offered its application as a SaaS solution, senior management determined that a prolonged outage could be disastrous for FMC and they initiated the push onto cloud infrastructure to mitigate the risk of outages.

When FMC moved to the cloud, it contracted with its hosting provider to periodically copy an image of its virtualized system images and replicate these to a secondary location. It also periodically replicated the customer data so it would be available whenever an event that threatened service continuity occurred. Working with each of its customer, FMC established recovery point objectives (RPO) and recovery time objectives (RTO) and defined these as part of their hosting provider agreement. This approach was less expensive to FMC since it no longer needed to have a QA infrastructure capable of supporting the full production load, and FMC could quickly turn up a copy of the production environment at the alternate site as needed. Although the level of effort for planning the recovery didn't change, the costs and risks associated with implementing a full recovery plan were more manageable.

Service Level Management

Service Level Management defines the level of service needed to support customers. Once a company defines this service level for an application, it needs to monitor that application to ensure that service delivery objectives are met. The formal definition of a service level can be documented in an SLA and used to manage all vendors participating in the service delivery process.

In its early days, FMC established an application availability objective for its product line. When it moved to the ASP delivery model, its customers pushed FMC to expand on this with service level objectives for application availability and transaction response times. When FMC reengineered its application as a SaaS solution, it recognized the need to license an end-user experience monitoring package that allowed it to measure the actual transaction response times of a customer's end-users. Although the monitoring package was an expensive solution, FMC was able to justify the cost by avoiding expensive service level penalties.

When FMC moved to the cloud infrastructure, it tried to find a hosting provider offering the monitoring software as part of the hosted solution. It found one that included the capability for an inexpensive monthly fee rather than a full product license. This saved FMC additional licensing fees. In addition, the monitoring product helped FMC with problem management by pinpointing the root cause of performance-related incidents as they occurred. This helped FMC improve the quality of its software while also avoiding expensive service-related penalties.

WHITE PAPER: Cloud Computing

Summary

Although the use of key ITIL processes by FMC was hypothetical, the evolution of its process development over the past 25 years is similar to those of many of Savvis' customers. In particular, the evolution to a cloud computing platform can deliver additional benefits gained from provisioning automation. As we've seen, standardized builds simplify IT support processes and provisioning automation makes it even easier to support complex application architectures.

Less human interaction in basic infrastructure component builds represents a stabilizing influence to production infrastructures. The seven ITIL processes' benefits are as follows:

Change Management, Problem Management, Release Management	Increases knowledge about a platform being changed; decreases types of problems due to variability introduced during manual builds; improves quality testing of software releases
Security Management	Can automatically patch systems to keep infrastructure software at current recommended levels
Capacity Management	Can quickly deploy systems for new application development, QA testing, and production workload changes
Service Continuity Management	Reduces costs of maintaining a production-ready recovery site with no impact on project schedules due to QA environment reuse
Service Level Management	Improves overall infrastructure availability as a by-product of improved Change, Problem, and Release Management

About Savvis

Savvis, Inc. (NASDAQ:SVVS) is an outsourcing provider of managed computing and network infrastructure for IT applications. By outsourcing to Savvis, enterprises can focus on their core business while Savvis ensures the quality of their IT infrastructure. Leading IT organizations around the world have selected Savvis to help them improve their service levels, reduce capital expense, and deal with the rising costs of bandwidth, energy, real estate, staff, and expertise. As a pioneer in utility computing, Savvis understands and harnesses the latest advances in technology like virtualization, cloud computing, and support process automation.

For more information about Savvis, visit www.savvis.net or call 1.800.SAVVIS.1 (1.800.728.8471).

EMEA
Savvis UK Limited
Tel +44 (0)118 322 6000

ASIA PACIFIC
Savvis Singapore
Company Pte Ltd
Tel +65 6768 8000

JAPAN
Savvis Communications K.K.
Tel +81.3.5214.0151