

JANUARY 2014



Mark Allen Miller

CORPORATE FINANCE PRACTICE

The enterprise IT infrastructure agenda for 2014

IT infrastructure managers must simultaneously capture the next rounds of efficiencies, accelerate the transition to next-generation infrastructure, reduce risks, and improve organizational execution.

**Björn Münstermann,
Brent Smolinski,
and Kara Sprague**

This year has been tough for many organizations that manage IT infrastructure—the hardware, software, and operational support required to provide application hosting, network, and end-user services. Highly uncertain business conditions have resulted in tighter budgets. Many infrastructure managers have rushed to put tactical cost reductions in place—canceling projects, rationalizing contractors, extracting vendor concessions, and deferring investments to upgrade hardware and software.

We have conducted more than 50 discussions with heads of infrastructure at *Fortune* Global 500 companies over the past six months to get a sense of the issues they are wrestling with. Clearly, infrastructure leaders must

meet 2013 budgets while ensuring they can address critical challenges in 2014 and beyond. They can do so by pulling 11 levers.

Capture the next round of efficiencies

There is no indication that 2014 will be dramatically easier from a budgetary standpoint than 2013 has been at many companies. Even as infrastructure organizations lock in 2013 savings, they need to take several actions to establish a pipeline of cost-improvement initiatives that will create space in their budgets for 2014 and 2015.

Takeaways

In the face of continuing uncertainty that's resulted in tighter budgets, IT infrastructure organizations in 2013 have had to move quickly to make tactical cost reductions.

Next year, budgetary pressures are unlikely to let up—but leaders can take on critical challenges by employing 11 levers.

Realizing new efficiencies to free up resources, moving quickly toward more flexible and scalable next-generation infrastructure, mitigating risks in alignment with business priorities, and improving execution for seamless delivery will be key.

1. Put in place a commercial-style interaction model with business partners

Traditionally, business partners and application-development managers have argued that they do not understand and cannot influence large infrastructure expenditures, which complicates demand management. In response, large infrastructure functions are establishing commercial models for interacting with their business partners. These models involve several efforts:

- continuing to implement standard service offerings that can be consumed on a price-times-quantity basis
- creating bottom-up unit costs for each service based on a detailed bill of materials
- investing in integrated tools to automate the data collection, aggregation, analysis, and reporting required for cost transparency
- putting in place the roles required to interact with business partners in a more commercial way—including product managers who can define standard offerings and solutions architects who can help developers combine the right mixture of standard offerings to meet a business need

2. Use project teardowns to optimize the cost of new demand

Solutioning—the process of converting a set of business requirements into specifications that can be implemented—is one of the biggest drivers of infrastructure costs. Fairly granular decisions about which operating-system version, server model, and storage tier to use can affect the cost to host an application by a factor of ten. Just as

consumer-electronics companies disassemble or “tear down” their own products and those of competitors, infrastructure organizations must apply similar thinking to new projects and existing systems.

This involves a structured process to lay out the full set of options in hosting an application; mapping the dependencies among decisions (for example, among different layers in the stack); assessing the cost, performance, and risk implications of each decision; and engaging business partners on important trade-offs. This process can be applied first to new projects and then extended to large existing applications over time.

3. Build industrial-style procurement capabilities

Procurement of hardware, software, and services required to operate an enterprise environment is becoming more challenging for senior infrastructure managers. Even as more procurement spending is devoted to software, many infrastructure organizations continue to use techniques developed for hardware procurement. These techniques are not entirely effective given software's product fragmentation and relatively high switching costs. Legacy contracts, sometimes with unrealistic volume or revenue commitments,¹ create financial and strategic constraints as infrastructure managers try to reshape their organizations and budgets to meet revised business expectations for efficiency and flexibility. To add even more complexity to managing vendor relationships, large vendors are starting to make incursions across technology domains (for instance, network-equipment vendors are providing servers, and server vendors are providing data storage). Despite the added complexity, though, these developments increase competition and provide additional opportunities.

¹A contract with a software vendor usually requires the customer to pay a certain amount even if changing organizational requirements mean it no longer needs the same level of usage.

Even after years of consolidation, most infrastructure leaders work in environments that they believe are too inflexible.

In response, infrastructure organizations should adopt purchasing techniques developed by automobile and consumer-electronics manufacturers, which source billions of dollars of components each year. Techniques that infrastructure organizations should employ include the following:

- creating tight integration between product-design and procurement decisions, building cross-functional teams, and rotating managers between procurement and product-management roles
- investing in procurement as a core discipline, separating strategic category-management roles from transactional purchasing-execution roles, and ensuring procurement staff have required business and product knowledge
- developing an information advantage by comparing vendors, thoroughly analyzing vendor costs, and investing in accurate inventory management
- building credibility when negotiating with vendors by making sure there is a single lead for all negotiations and using signaling mechanisms such as press coverage to indicate willingness to switch vendors
- employing differentiated strategies to achieve the lowest possible cost by using varied procurement tactics (depending on what is appropriate given the market structure) such as sole source, auctions, and long-term contracts, as well as conducting internal discussions about the lowest price vendors might accept and how far the company is prepared to push each vendor

- ensuring that the company's and the vendors' incentives are aligned by using open-book pricing, making year-on-year cost reductions, and mutually sharing the gains

Accelerate the transition to next-generation infrastructure

Even after years of consolidation and standardization, which have led to huge improvements in efficiency and reliability, most infrastructure leaders work in environments that they believe are too inflexible, provide too few capabilities to business partners, and require too much manual effort to support. Addressing these problems to create more scalable and flexible next-generation infrastructure will require sustained actions in multiple dimensions.

4. Determine how to influence application-development road maps to enable more scalable and efficient infrastructure

Most large institutions have programs under way to develop and roll out private-cloud environments in order to reduce hosting costs and dramatically improve the speed of delivery. Adopting techniques pioneered by “hyperscale” infrastructure functions serving e-commerce and social-media organizations can significantly enhance the business case for next-generation hosting environments.

These techniques include extensive self-service or automation, software-defined networking, commodity components, and aggressive use of open-source technologies. They have enabled some companies to reduce hosting costs by 50 to 75 percent.

However, infrastructure functions must take into account several important considerations as they evaluate how radically they can evolve their hosting environment:

- What will it take for application developers to learn how to use self-service capabilities effectively?
- Given existing architectures, should legacy applications be migrated to the new environment, or should it primarily be used for new applications?
- What are the performance implications of commodity components, given the critical workloads?
- What technical skills are required to build and support an environment that leverages hyperscale technologies?

5. Get the operating model in place to scale private-cloud environments

Many large infrastructure functions are experiencing “cloud stall.” They have built an intriguing set of technology capabilities but are using it to host only a small fraction of their workloads. It may be that they cannot make the business case work due to migration costs, or that they have doubts about the new environment’s ability to support critical workloads, or that they cannot reconcile the cloud environment with existing sourcing arrangements. Over the next year, infrastructure organizations must shift from treating the private cloud as a technology innovation to treating it as an opportunity to evolve their operating model. This involves a number of elements:

- tightly integrating private-cloud offerings into their service catalog and establishing business rules to make deployment of these offerings the default option for many types of workloads

- concentrating responsibility for the private-cloud offering to a single owner that oversees its economics and service-level performance
- redesigning operational processes to eliminate manual steps required for traditional environments
- leveraging automation to facilitate DevOps² and give developers more control over their applications (within guidelines)
- recasting sourcing arrangements to enable cloud operating models³; many infrastructure organizations are trying to migrate from traditional sourcing arrangements to virtual private-cloud models, and others are seeking models in which they control the developer interface and a provisioning or orchestration layer but may source the underlying servers in an integrated way

6. Advance end-user offerings to facilitate business productivity

After all the attention paid to application hosting over the past several years, many infrastructure leaders have started to conclude that they need to increase the attention and focus they devote to innovating end-user capabilities. At many companies, the most critical employees (in functions such as sales, marketing, research, and design) depend heavily on end-user technology tools such as e-mail and calendar rather than on business applications such as customer relationship management to enhance their productivity.

However, there is a fair degree of uncertainty about the ultimate direction of end-user capabilities. Infrastructure leaders are looking for answers to the following questions:

- How do we strike the right balance between security and mobility, that is, the use of

²DevOps is a software-development method that stresses collaboration, communication, and integration among software developers and IT professionals.

³See James Kaplan, Chris Rezek, and Kara Sprague, “Protecting information in the cloud,” January 2013, mckinsey.com.

smartphones, tablets, and other cloud-enabled devices that extend the reach of the company's wired information infrastructure but also make the information more vulnerable to breaches?

- Where and how should we deploy virtual-desktop⁴ infrastructure?
- Is there a practical business case for having unified communications⁵ and rich collaboration capabilities?
- Does the productivity impact of desktop videoconferencing justify increased bandwidth costs?

Answering these questions will require deep engagement not only with business-unit IT functions but also with business managers and frontline personnel, who can help leaders develop a granular understanding of frustrations with existing tools and gain insight into how more sophisticated ones might integrate with day-to-day business operations. In doing this, infrastructure managers should pay particular attention to integration across tools. There is typically as much of an opportunity in tightening the linkages across existing capabilities as there is in adding entirely new functionality.

Reduce risks

As more business value migrates online and business processes become more and more digitized, IT infrastructure inevitably becomes a bigger source of business risk. Customer-facing systems could slow to a crawl because of insufficient computing capacity, data-center outages can disrupt business, and critical intellectual property could be

extracted from inadequately secured networks. It is easy to overreact and select strategies that reduce risk but carry high costs with regard to capital expenditures or reduced business flexibility.

To serve their business partners effectively, senior infrastructure managers will have to create easily understood options that allow business partners to make practical trade-offs between cost and risk.

7. Ensure facility footprints provide required resiliency at acceptable cost

Hurricane Sandy, which came little more than a year after the Japanese tsunami, was a sobering moment for many infrastructure organizations on the east coast of the United States. A few institutions suffered severe outages; many more had close calls.

The experience accelerated the years-long process of companies moving servers and other assets out of closets and other subfunctional facilities into consolidated, strategic data centers.

It also reinvigorated a long-running debate about the pros and cons of real-time failover⁶ versus geographic diversity. Is it better to build data centers in close-together pairs so applications can run synchronously across the two and so avoid downtime if one facility is impaired—but still run the risk that an extreme event such as a tsunami or an earthquake might impair both facilities? Or is it better to accept a small period of downtime so that, in the event of a disaster, applications can be brought back up in a facility hundreds of miles away? Might it be better to expend the capital so that applications run synchronously across a data-center pair, with recovery to a third, remote facility if required?

⁴Virtual-desktop services separate the end-user "desktop" environment from a physical machine and can potentially simplify support, improve data security, and allow users to access applications and data from a range of devices.

⁵Unified-communications platforms integrate chat, e-mail, knowledge management, video, and voice for a seamless user experience.

⁶Real-time failover refers to the ability to move from your primary capabilities to disaster-recovery capabilities in real time.

Naturally, different organizations will have different answers to these questions, but there are a few musts in addressing the issue:

- starting with business applications and processes and being willing to create segments to avoid “leveling up” to the most expensive answer for all applications
- integrating modular and predesigned architectures into data-center-build plans to increase flexibility and realize lower costs for an appropriate level of resiliency
- looking at resiliency across the entire stack—more robust facilities may not be the lowest-cost or most effective mechanism for increasing the resiliency of a set of applications

8. Instrument the environment to support next-generation cybersecurity

Every week seems to bring news of another cyberattack, motivated by financial gain, political activism, or—in some cases—national advantage. In many companies, the policy aspects of cybersecurity are being moved out of infrastructure organizations in order to enhance their visibility and proximity to other functions like enterprise-risk management.

But infrastructure has an utterly critical role to play as next-generation cybersecurity strategies are put in place. Increasingly, cybersecurity will depend on sophisticated intelligence and analytics about attacker strategies. This requires massive amounts of data about what is transiting the enterprise network, where it is coming from, and who has been accessing critical systems and data.

However, infrastructure functions have to partner with security functions to make



effective trade-offs so they can find ways to extract massive amounts of granular data without compromising performance or creating too much additional complexity.

Improve organizational execution

Organizational capacity is one of the perennial frustrations in enterprise infrastructure—there is never enough managerial, operational, or technical talent to provide day-to-day service delivery while pursuing necessary improvements. Creating required organizational bandwidth will depend on infrastructure managers pulling organizational, performance-management, and talent-sourcing levers.

9. Make the transition to a plan-build-run organizational model

Traditionally, large infrastructure functions have been organized according to a combination of regions and “technology towers.”⁷ However, both of these models seem to be hitting their limits for large enterprises. Regional constructs constrain infrastructure functions’ ability to get the most from their investments in new technologies

⁷A technology tower is a grouping of related technologies, such as servers, storage, end-user technology, and network.

and make it harder to support global business processes and applications. Moreover, private-cloud environments, converged infrastructure products, Internet Protocol telephony, unified communications, and virtual-desktop infrastructure all make traditional distinctions among end users, data centers, and networks less and less relevant.

In response, leading infrastructure organizations are putting in place functional organizations with distinct “plan,” “build,” and “run” capabilities.

- “Plan” includes service or relationship management, which is responsible for collecting business requirements, performing demand management, and serving as the overall interface with business partners. It also includes product management, which is responsible for developing and optimizing a set of reusable service offerings to be consumed by business partners.
- “Build” includes product engineering and deployment. Product engineering designs and configures the technology to make the service offerings defined by product management ready for use in a production environment. Deployment takes requests from service management, develops implementable solutions using standard service offerings, and provisions them into the day-to-day environment.
- “Run” performs all the operations and support to keep the technology environment running and meeting service-level expectations.

10. Proactively create the next generation of infrastructure business leaders

There is a long-standing model of career progression within most infrastructure organizations. Junior engineers acquire technical

expertise in a given area and become more and more specialized in storage, databases, or networks. Frontline operations managers demonstrate their ability to keep things running and manage larger and larger operational teams over time.

This model creates technology specialists and effective operators. It does not necessarily create business leaders who can drive innovation or technology integrators who can solve problems that span the server, storage, network, database, and middleware domains.

To expand the pool of effective infrastructure managers—and so expand the organization’s ability to do big things—senior infrastructure leaders will have to broaden their set of talent-management levers by doing several things:

- increasing their use of rotational staffing—moving high performers across technology domains as their careers progress
- expanding the hiring aperture by interviewing and hiring managers from both application-development and commercial-technology vendors
- creating nontechnical training to help managers build skills in general technology problem solving and develop knowledge of the businesses they support

11. Drive performance management to the front line

There are huge variations in productivity and quality (typically by a factor of four to ten) between the most and least effective help-desk agents, system administrators, database administrators, desktop technicians, and other frontline infrastructure staff.

While most infrastructure functions have thick books of metrics—such as help-desk first-call resolution, outages by severity, number of scheduled job failures, and mainframe utilization—almost all of these metrics measure platforms or parts of the organization, not the performance of frontline personnel. As a result, weaker performers do not get the coaching they need, strong performers do not get the recognition they deserve, and organizational productivity and quality suffers.

To continue making advances in efficiency and quality and to meet business expectations, infrastructure organizations will have to not only implement frontline metrics, such as tickets closed per day, but also use them in performance huddles and one-on-one coaching to improve individual performance.

In many cases, implementing these changes may require reversing work-from-home policies and bringing frontline staff back into operational centers where they can engage with managers and collaborate with peers.



We hope the practices described in this article will help infrastructure managers navigate a challenging, and sometimes conflicting, set of demands to meet budget constraints, protect the business, and provide innovative capabilities in 2014 and beyond. Naturally, no single infrastructure function will address all of these areas simultaneously—how and when they can will be determined by a combination of the starting point, business needs, and organizational constraints. ○

The authors wish to thank James Kaplan for his contributions to this article.

Björn Münstermann is a principal in McKinsey's Munich office, **Brent Smolinski** is an associate principal in the Atlanta office, and **Kara Sprague** is a principal in the San Francisco office. Copyright © 2014 McKinsey & Company. All rights reserved.