AVASANT

Data Center Management

In a Shared, Multi-Customer Environment

November 2010

This paper is intended to provide an overview of the fundamental elements of data center management in the context of a multi-customer environment, which is generally characteristic of outsource data center providers, but also increasingly similar to in-sourced data centers providing services to multiple internal customers in a best practice, shared services organization model.

Introduction

The management of today's data centers is becoming increasingly complex. As data center managers are charged with optimizing resources and minimizing capital-intensive upgrades, at the same time they are grappling with how best to incorporate new technologies like cloud computing, virtualization and new data storage applications. Among other forces data center managers need to address include the drive to lower the total cost of ownership (TCO), while increasing quality of service; energy provisioning and consumption in a "green IT" era; increasingly stringent information security requirements; data center design efficiency; business continuity; and the IT Infrastructure Library (ITIL).

- This paper is intended to provide an overview of the fundamental elements of data center management in the context of a multi-customer environment, which is generally characteristic of outsource data center providers, but also increasingly similar to in-sourced data centers providing services to multiple internal customers in a best practice, shared services organization model. The topics addressed include:
- The typical scope of data center services provided to customers, associated tasks and activities, and key ITIL processes to institute;
- The key roles required for the successful delivery of data center services, a recommended organizational concept and key joint governance mechanisms required to help govern customer-provider agreements;
- The typical Service Level Agreement (SLA) measurements that drive levels of performance to be managed to in order to meet customer needs;
- The customer pricing/chargeback mechanisms used to recover data center costs of providing services; and
- The key challenges in managing data center operations and some solution approaches to address them.

Facility Considerations

As the data center has evolved from the monolithic "glass house" of the mid-1900's to today's mix of complex, heterogeneous computing platforms, the design of the facility itself requires joint expertise in both facilities design and computing environment design. New data centers need to be designed and built in as modular a fashion as possible, allowing flexibility and scalability to accommodate both growth and likely increases in power and cooling energy needs, the cost of which is fast approaching the number one single cost of operating a data center. The importance of these considerations in the planning for building new data centers is illustrated by the ageing of older data centers, many of which cannot economically be refurbished to meet current needs, as their design did not anticipate today's requirements. As a rule of thumb, data centers require major refurbishment every seven years in order to accommodate growth and change. Costs for periodic refurbishment should be incorporated into the business case for building or procuring a data center in order to properly reflect the total cost of ownership.

Data Center Classification

When considering building or refurbishing a data center, service providers need to determine the applications availability requirements of their customers. There are large differences in the availability requirements of applications used to support worldwide equity trading, for example, and those used to support less time sensitive administrative functions. The higher the availability requirements, the more capability and redundancy need to be designed into the data center. To help determine data center design based

on availability requirements, the Uptime Institute has developed a data center tier approach, which has been accepted as the standard in data center availability classification. The simplest is a Tier 1 data center, which is basically a server room, following basic guidelines for the installation of computer systems. The most stringent level is a Tier 4 data center, which is designed to host mission critical systems, with fully redundant subsystems and compartmentalized security zones controlled by biometric access controls methods. The chart below displays the four tiers and corresponding, high level requirements.

Tier Level	Requirements
1	• Single non-redundant distribution path serving the IT equipment • Non-redundant capacity components • Basic site infrastructure guaranteeing 99.671% availability
2	• Fulfills all Tier 1 requirements • Redundant site infrastructure capacity components guaranteeing 99.741% availability
3	• Fulfills all Tier 1 & Tier 2 requirements • Multiple independent distribution paths serving the IT equipment • All IT equipment must be dual-powered and fully compatible with the topology of a site's architecture • Concurrently maintainable site infrastructure guaranteeing 99.982% availability
4	• Fulfills all Tier 1, Tier 2 and Tier 3 requirements • All cooling equipment is independently dual-powered, including chillers and Heating, Ventilating and Air Conditioning (HVAC) systems • Fault tolerant site infrastructure with electrical power storage and distribution

Data Center Outsourcing

Historically, the management of data center services has been outsourced to a single provider along with other infrastructure services, such as data/voice networks, desktop support and help desk. Sometimes, applications services were also outsourced to the same provider. As customers experienced service inconsistencies across service areas, long lock in periods, inability to meet financial deal objectives and relationship issues with a single provider, the trend began to move towards more selective outsourcing of service areas and shorter multi-year agreements (currently 3 to 5 years is common today).

Another important trend emerging in the last 3 years is the concept of Remote Infrastructure Management or RIM, in which most of the outsourced data center staff is located offshore. Enabled by the advent of more robust and secure communications capability and remote management tools, this concept of remote management provides significant savings compared to the traditional on-site staff model.

Providers now offer a flexible set of RIM services, enabling the outsourcing of various data center services without requiring customers to move equipment to the providers' facilities - an 'asset-light' model that does not require large, upfront capital investment. In addition, customers can take a more incremental approach by outsourcing monitoring and first level support, while retaining the management of the facility, equipment and higher level support internally.

The Forrester chart below illustrates the types of infrastructure outsourcing and the level of maturity for each in the graphic below. Note that data center outsourcing is reaching maturity, while offshore infrastructure management is in the emerging stage.

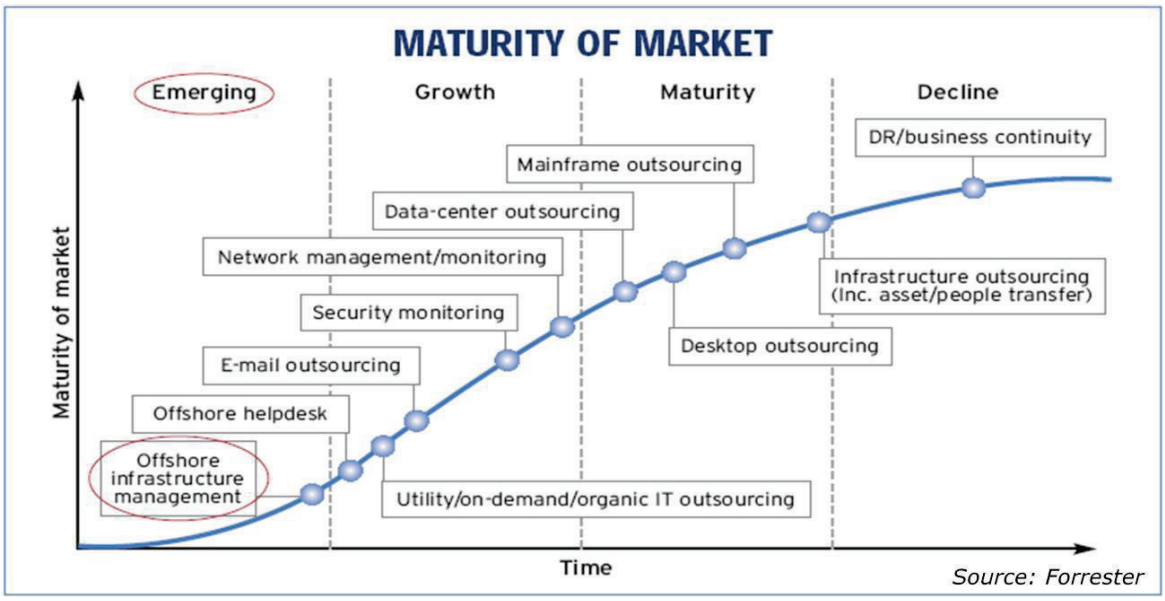

Data Center Services

Typical data center services provided and managed by data center services providers are described below.

Key Data Center Services
Server Operations and Administration - the provisioning and day-to-day management of the Data Center server environment, providing a stable server and supporting infrastructure. Activities include monitoring, job scheduling and operations, file/data transfer, system administration and server management.
Storage and Data Management Services - the provisioning and day-to-day management of the data storage environment and providing a stable supporting infrastructure. Activities include monitoring performance, optimizing utilization and performing data backup and restore functions.
Remote Access Services - the installation, management, operations, administration and support of the hardware and software that enable secure Remote Access to all authorized systems (e.g. Citrix Metaframe via dial up and Internet, Web-based e-mail, terminal services, etc.). Activities include in-

stalling, testing, configuring and providing technical support for RAS hardware and software.
Applications Support Services – the provisioning of the infrastructure and operational activities associated with the installation, operations, administration and support of designated Customer Applications (e.g., business applications, Web applications). Activities include supporting applications test-to-production activities, installing applications software and performing applications software distribution.
Database Administration Services - the installation, management, maintenance and support of customer database solutions and configurations (e.g. stand-alone databases, clustered databases). Activities include data set placement, database performance monitoring, data recovery, and data security and integrity maintenance at the physical level.
Middleware Administration Services - the maintenance and support of middleware that enables integration, EDI, FTP and system interfaces between customer applications and third party applications. Activities include implementing configuration change, upgrade and refresh requests, load balancing and tuning.
Messaging Services - the provision and management of a messaging platform (e.g., Microsoft Exchange) and messaging services used to directly connect customer users to these services via Local Area Network (LAN) or the Internet, including connectivity with desktops, laptops, handheld devices (e.g., BlackBerry) and fax-enabled devices (e.g., RightFax). Activities include administering email services and implementing, configuring, maintaining and managing the messaging services infrastructure.
End User Administration Services - the management and coordination of account activation, termination, changes, expirations and end user support resources. Activities include administering User ID and system access consistent with customer information security requirements, performing password resets and creating system access procedures for new customer applications.
Web Support Services - the provisioning of web infrastructure and middleware support for customer's web applications, including internet and intranet portals, internal collaboration tools and social media. Activities include installing, configuring and supporting web infrastructure components, managing internet bandwidth and web proxy servers, and monitoring and reporting on web site activity.
Data Center Security Services – the management and maintenance of the physical and logical security of all Data Center Services components (e.g., hardware, software) and data, malware protection, access protection and other data center security services in compliance with customer security requirements and all applicable legal and regulatory requirements. Activities include implementing a security incident response program, maintaining controlled access to computing areas, disabling access by terminated user staff, monitoring unauthorized access attempts, deploying anti-malware updates and patches and promptly notifying customer of security breaches.
Technology Refreshment and Replenishment (TR&R) – the modernization of the IT infrastructure on a continual lifecycle basis to ensure that the system components stay current with evolving industry standard technology platforms. Activities include developing a TR&R plan for each in-scope service area, implementing the TR&R plan and periodically reporting TR&R plan progress.
Maintenance Services – the maintenance and repair of hardware, software and networks to include break/fix and software maintenance. Activities include deploying software and firmware updates, patches and hot packs; performing major software release upgrades; maintaining software release matrices across development, QA and production environments; and replacing defective hardware components.
Integration and Testing Services – the testing required to ensure that all individual data center components (e.g., hardware, software, middleware, interfaces and network) configured with or added to the environment work together cohesively to achieve the intended results and meet customer requirements. Activities include developing unit, system, integration and regression test plans; conducting integration and security testing; and updating the Configuration Management Database (CMDB) for testing-related changes.
Asset Management Services - the ongoing management and tracking of new and upgraded data center components (e.g., hardware, software) in the customer's asset management system. Activi-

ties include receiving, inventorying, installing, applying preventive maintenance to and disposing of assets.
Software License Management Services - the acquisition, ongoing management and tracking of in-scope software licenses. Activities include identifying and reporting license compliance issues, performing license audits and reconciling number of licenses to the number of installs.
IT Service Continuity and Disaster Recovery (DR) Services - the provision of contingency plans to ensure continuous availability of customer application functionality under all circumstances, including recovery in the event of a disaster. Activities include providing a remote site enabling the recovery of critical applications, developing/maintaining a DR plan, performing DR testing and review, and coordinating DR recovery procedures in response to a disaster.

As advances in server and data storage technology have been introduced, so too have advances in software management tools been made in response to customers' needs for increased availability and fast resolution when problems occur. Tools for monitoring operations, for example, can be used to detect potential problems before they occur and issue alerts to the operations staff.

Along with these advances, the ways in which data center services are delivered from a "Best Practice" perspective have evolved. The Information Technology Infrastructure Library (or ITIL) is a set of concepts and practices for IT Services Management (ITSM), Information Technology Development and IT Operations developed by the UK Office of Government Commerce. ITIL gives detailed descriptions of a number of important IT practices and provides comprehensive checklists, tasks and procedures that any IT organization can tailor to its needs. The ITIL process model for IT infrastructure management and operations defines best-practice processes for IT service delivery and support. ITIL is prescriptive about the tasks involved in those processes and promotes a common language for IT infrastructure services. As such, ITIL has been adopted as an IT industry standard for managing and delivering IT infrastructure operational services and its compliance has become a de-facto requirement for all IT infrastructure service providers. The key ITIL processes related to IT infrastructure services are described in the table below.

Key ITIL Processes
Service Level Management – the process of defining, agreeing, documenting and managing the levels of customer IT service required to ensure the service levels are delivered as agreed in the SLAs in a cost effective, secure and efficient manner.
IT Financial Management – the discipline of ensuring the IT infrastructure is obtained at the most effective price and calculating the cost of providing IT services so that an organization can understand the cost of its IT services. These costs may then be recovered from the customer of the service.
Capacity Management - the process of matching IT resources to business demands. Activities include application sizing, workload management, demand management, modeling, capacity planning, resource management and performance management.
Continuity Management - the process of proactively ensuring plans are put in place and managed to ensure IT services can recover and continue even after a serious incident occurs. Activities include conducting a Business Impact Analysis (BIA), performing risk analyses, evaluating options for recovery, producing a contingency plan and testing the plan on a regular basis.
Availability Management – the process of allowing organizations to sustain the IT service availability to support the business by ensuring each IT component performs at an agreed level over a period of time. Activities include identifying availability requirements, compiling an availability plan, monitoring availability and monitoring maintenance obligations.
Incident Management – the process of restoring normal operations as quickly as possible with the least possible adverse impact on the business, thus ensuring that the best possible levels of service quality and availability are maintained, consistent with SLAs.
Problem Management – the process of resolving problems that causes incidents in order to minimize the impact of incidents on the business and to prevent their recurrence. Activities include trend

analysis and root cause analysis, comprising problem identification, classification, investigation and diagnosis.
Change Management – the process of introducing a change to configuration items with minimal disruption of services, a reduction in back-out activities and economic utilization of resources involved. Activities center on developing standardized methods and procedures for efficient handling of all changes.
Configuration Management – the process that tracks all individual Configuration Items (CIs) in a configuration management database (CMDB) in order to accurately track infrastructure components.
Release Management – the process of controlling the introduction of new or changed software and/or hardware into the IT infrastructure to protect the live environment by ensuring the availability of licensed, tested and version-certified software and hardware, which functions as intended. Activities include planning the implementation, designing and implementing procedures for IT changes, managing expectations of customers, and controlling the distribution and installation of changes.
Service/Help Desk – the service support discipline focused on the user of IT services and primarily concerned with ensuring they have access to the appropriate services to support business functions. The service desk functions as a single-point-of-contact for end-users' incidents and owns the incident management process end-to-end, keeping customers informed of status until incidents are resolved.

Data Center Services Key Roles

In order to deliver services effectively, data center services providers develop key roles derived from the services and ITIL processes described in the previous section. Responsibilities and accountabilities are assigned to these key roles, which are then staffed at the required levels with appropriately-skilled people.

Typical data center management roles include:

Key Data Center Provider Roles	
1. Customer Account Manager	• Builds Relationship with Customer IT Management • Facilitates IT Initiatives • Helps customer develop business case for initiatives with data center investment impact • Responsible for Customer Satisfaction
2. Service Delivery Managers	• Manages day-to-day operational services for assigned areas such as server and network operations, project support, service desk, data administration, data storage, systems administration • Addresses service delivery issues/problems • Faces off to customer IT • Publishes regular operational status reports
3. Performance Manager	• Tracks and reports on provider's compliance with Service Level Requirements

	• Facilitates ITIL Compliance, Continuous Improvement & Benchmarking – Data Center
4. Finance Manager	• Develops pricing for customer agreements • Tracks and reports on actual charges vs. budgets • Works with customer finance representatives to facilitate monthly billing and payment • Assists in the development of business cases for projects
5. Procurement & Asset Mgt. Manager	• Procures and manages data center assets Monitors changes and updates asset Inventory • Maintains Contract and Resolve Contract Disputes • Manages License Compliance Verification
6. Contract Manager	• Manages contract negotiations • Maintains contracts • Resolves contract disputes
7. Security, Compliance & Risk Mgt. Mgr.	• Monitors Regulatory Compliance • Monitors Standards Compliance • Monitors Security/Controls Compliance • Manages Risk Reviews
8. Architecture Manager	• Develops and maintains data center architecture plan • Tracks new technologies Monitors Compliance with technical standards
9. Facilities Manager	• Maintains the data center facility in accordance with the targeted Tier Level (usually Tier 3 or 4) • Plans and executes major refurbishment (usually every 7 years)
10. Process Manager	• Provides ITIL process requirements to service delivery areas • Monitors ITIL compliance • Facilitates ITIL implementation in customer organization, when required

In addition to the key management roles required, technically-skilled staff are required in the following areas:
- Server platform operations and maintenance
- Network operations and maintenance

- Data storage, backup and recovery
- Service desk
- Innovation using new technology application (e.g., server consolidation, virtualization, data de-duplication, cloud computing, etc.)
- Physical and logical security.

Organization Structure

The most common structures for organizing these key management and technical roles is using a matrix consisting of service delivery areas (sometimes referred to as service towers) on one dimension and functions such as Process Compliance matrixes across service areas on the other. An example is depicted below.

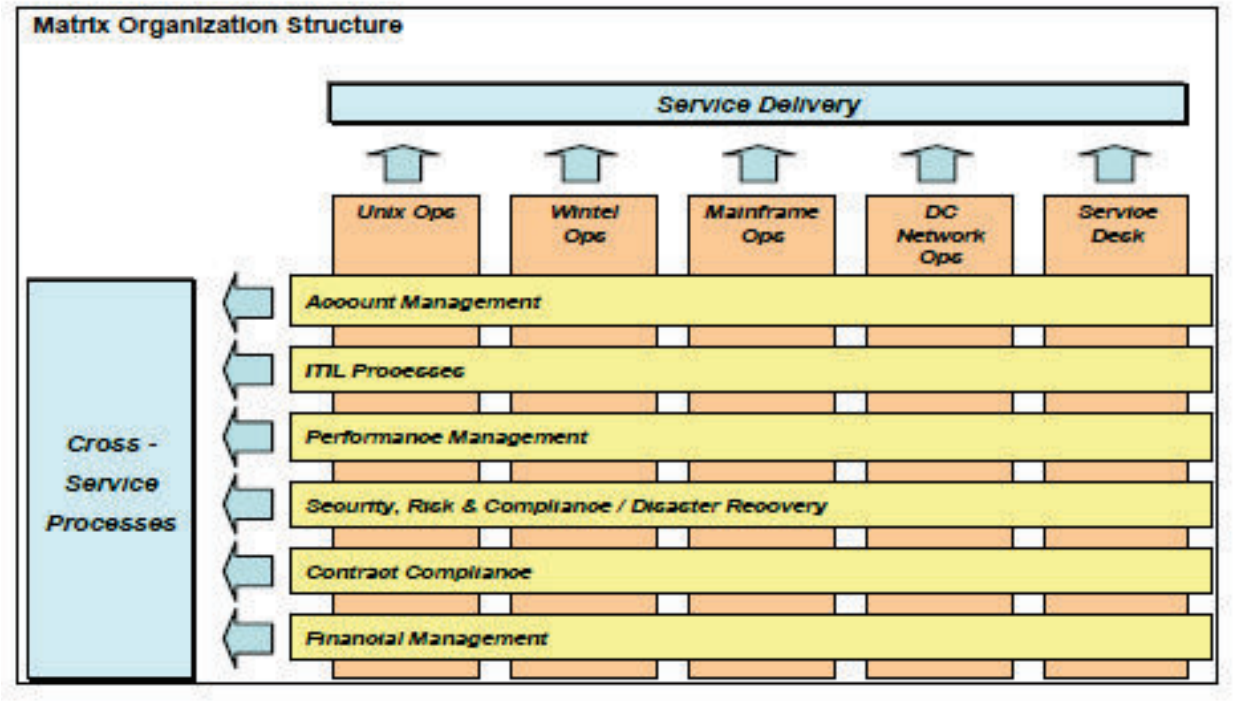

Governance Structure

In addition to the key management and technical roles described above, an important determinant of success is the institution of supporting governance bodies or committees to jointly govern each customer relationship and contract. The customer-provider relationship needs to be nurtured at the executive, management and operational levels in order for both parties to achieve their objectives of entering into a services agreement.

The governance process and supporting joint committees should be established during services transition and each should include a charter, fixed membership, standing agenda, regular meeting frequency, inputs and outputs. The operational level committee may meet weekly and focus on short term operational issues, while the management committee would meet monthly focusing on service performance and any issues escalated from the operational committee. In turn, the executive level committee may meet semi-

annually and focus on relationship quality and business alignment. The graphic below illustrates a best practice governance structure.

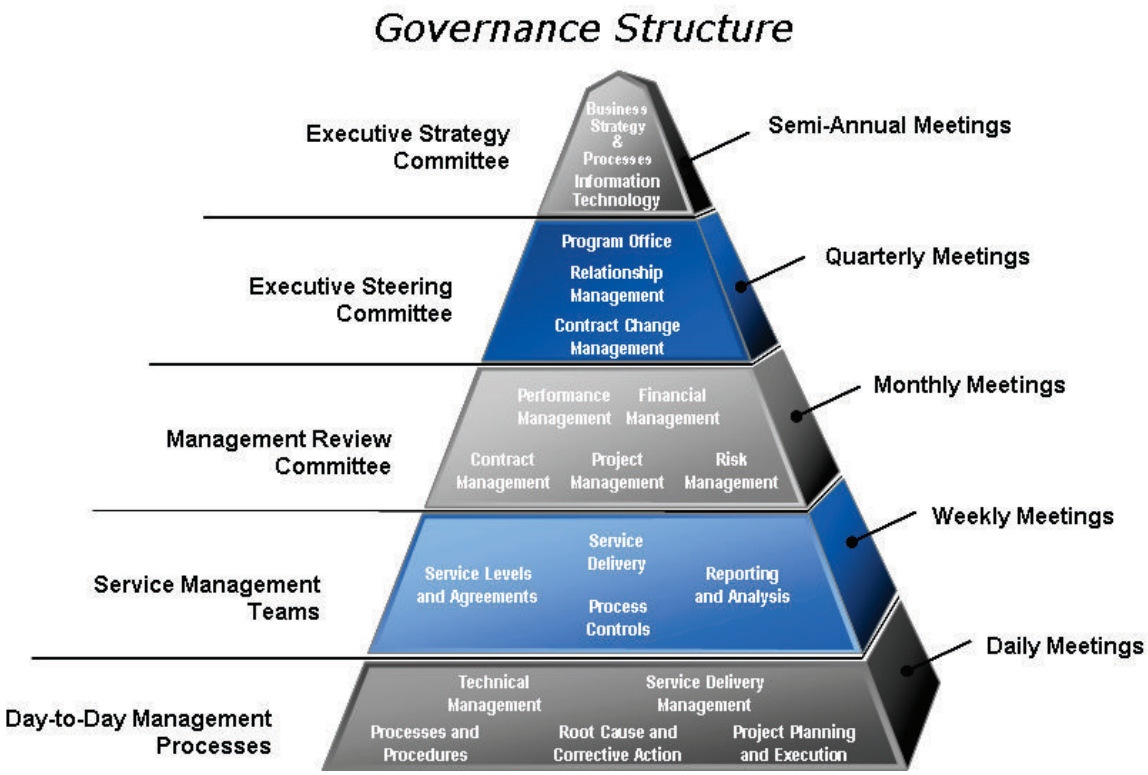

Data Center Service Level Agreements

Service Level Agreements (SLAs) are a critical part of any agreement between data center services providers and their customers, as they represent the measureable service performance required for customers to run their businesses. The two most meaningful types of SLAs have to do with availability (e.g., of a server platform or support staff) and response time (e.g., of a server platform or to resolve an incident). To illustrate system availability, most mission-critical applications are required to be available for use 99.99% of the time or more. Less critical applications require somewhat less availability, perhaps 99.7%, while other platforms used for non-operational applications, like development, require even less availability at 98%. Each of these different availability levels require different levels of hardware, software and staff resources to manage. The higher the availability required, the higher are the costs to achieve it.

SLAs are an important part of the negotiations leading up to an outsourcing agreement because they drive cost for the provider, who must ensure that the pricing will cover those costs with an adequate profit margin. In addition, penalties in the form of fee reductions are invoked when a provider fails to meet key SLAs. During the development of an outsourcing agreement, customers define the service levels required in order to meet the needs of their business operations.

Some of the typical SLAs that are included in data center outsourcing agreements are:

Data Center Services SLAs
System Availability - the availability of in scope infrastructure components required to conduct the normal business operation of customer application systems at full functionality, including server components (e.g., server CPU, memory, internal storage, Data Center LAN), external storage, system software and network connections. Measured as a percentage of time available per platform per month.
Unscheduled Downtime - Inability to conduct normal business operation due to unscheduled application downtime. Measured in number of events per customer application per time period.
Service Desk Notification - Notification of Priority 1, 2 or 3 Outages to customer within a stipulated time.
Batch Processing – Completion of core scheduled production batch jobs per customer's approved schedule as a percentage of batch jobs run. For on demand batch processing, time to initiation.
Administrative Functions - Response time to set up or modify Job Scheduler definition and dependencies; change authorized system access within a time period for a percentage of requests.
Storage Administration – Monitoring and reporting on storage capacity utilization and trends; response time on requests for allocation of additional storage capacity and administration requests.
Internet Web Application Deployments – Number of successful web deployments as a percentage of requested deployments in a time period.
Database Administration – Response time for instance creation and refresh, requests for User IDs, grants, revokes, table space creation and data definition requests; schema changes and stored procedures; and performance tuning and maintenance. Measured as a percentage of successful events divided by total events.
Virtual Server Provisioning – Response time to provision a virtual server from authorized request to completion. Measured as the percentage of successful events divided by total events.
Server Provisioning and Administration - Response time to provision a physical server from authorized request to completion and to fulfill server administration requests. Measured as a percentage of successful events divided by total events.
IT Continuity and DR - Time to recover customer application systems (as defined in the DR plan) after DR Incident has been determined. Measured by respective application and target recovery time.
DR Testing - Successful annual test completion of each customer application recovery within targeted recovery time.
Operational Reporting – Provide operational reporting within agreed time frames, including: • CPU Load Average • Disk Utilization • Memory Utilization • Daily Systems Management Report • Customer Application Availability and Performance Report • SLR Summary Exception Report.

While it can be tempting to create many SLAs, it is advisable to use only those measurable SLAs that truly reflect business requirements. This will serve to provide sharper provider focus, enable better manageability and provide more impactful fee reduction penalties. A well-established set of SLAs is important because it sets boundaries and expectations for performance-based contracts and becomes the vehicle for defining services, performance levels, accountability, consumption of resources and charging/pricing structure.

Data Center Services Pricing Methods

Common data center pricing models consist of a fixed management fee and variable fees based on services consumed. These variable fees are based on the concept of recovering variable costs (plus margin) through pricing of recognizable resource units (e.g., a type of server). The provider estimates cost pools in each resource unit-related service area and allocates the cost over the estimated number of resource units. For instance, a price per month to support a UNIX server is based on recovery of the costs incurred to manage that server (plus margin). In most data center outsourcing agreements, resource unit rates are fixed for each year of the agreement and typically reflect efficiency improvements through reduced rates in succeeding years.

Typical resource units include the following:

- Number of Wintel Servers by Size (small, medium and large, based on the number of processors)
- Number of UNIX Servers by Size (small, medium and large)
- Number of Virtual Servers
- Number of Virtual Operating System Instances
- Number of Terabytes of Storage Used
- Number of Database Instances
- Number of Mainframe MIPs

Each of these resource units is priced within a relevant range of volumes (commonly, plus or minus 5% of the estimated baseline quantity) within which the monthly pricing is unchanged. This practical approach allows for predictability, as both parties recognize that small fluctuations within 5% of the targeted volumes tend to balance out in the long run. When resource unit volumes do vary more than 5%, either more or less from the baseline quantity, then there are adjustments to the charge in the form of Additional Resource Units (ARCs) or Reduced Resource Units (RRCs). The rates for these ARCs and RRCs are pre-negotiated before the agreement is signed. In the case where resource unit volumes fluctuate more significantly than the ARC – RRC relevant range (typically 20% from the baseline volume), then rates are typically renegotiated.

Two other types of fees are common in outsourcing agreements. One is a fixed fee for disaster recovery services. The other is a time and materials hourly rate by staff skill type, which is used for unforeseen project work over and above day-to-day operational work. These hourly rates are also typically negotiated as a part of the executed agreement.

Data Center Management Challenges and Solution Approaches

Data Center Management Challenges

Organizations today face a number of data center management challenges. More demanding services are expected from a data center, including securing and protecting rapidly expanding volumes of data, supporting a growing number of mission-critical applications, managing highly complex and heterogene-

ous environments, meeting more challenging service level agreements (SLAs) and implementing a variety of emerging "green" business initiatives.

Compounding these challenges are the increasing costs of managing a growing data center infrastructure. According to analysts such as IDC in the graphic below, as the data center environment becomes more complex, management costs increase, diverting resources from the business.

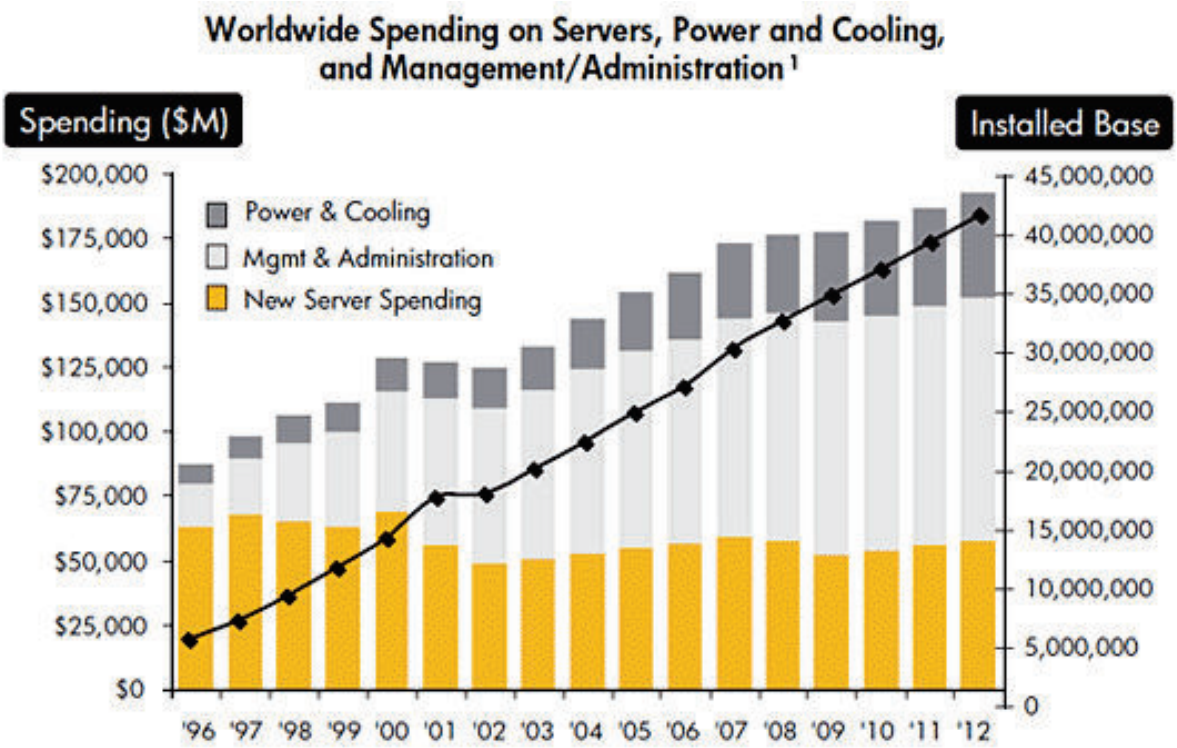

At Gartner's 2009 Data Center Conference, among the top ten most significant challenges faced by 654 survey respondents were in the following four areas:

- Data center space, power and/or cooling due to equipment sprawl (30+%)
- Developing a private/public cloud computing strategy (9%)
- Virtualization (8%)
- Finding/retaining IT talent (5%)

Data Center Facility Issues

By far, issues related to the data center facility (space, power and cooling) constitute the top challenges. Over the years, the density of computing equipment has been increasing steadily (i.e., smaller circuits, multiple cores, blade servers, etc.). A decade ago 500-1000 W per rack was typical, but currently a 42 U rack can easily draw 10-20 kW which can generate heat in the range of 30,000-70,000 British Thermal Units (BTUs).

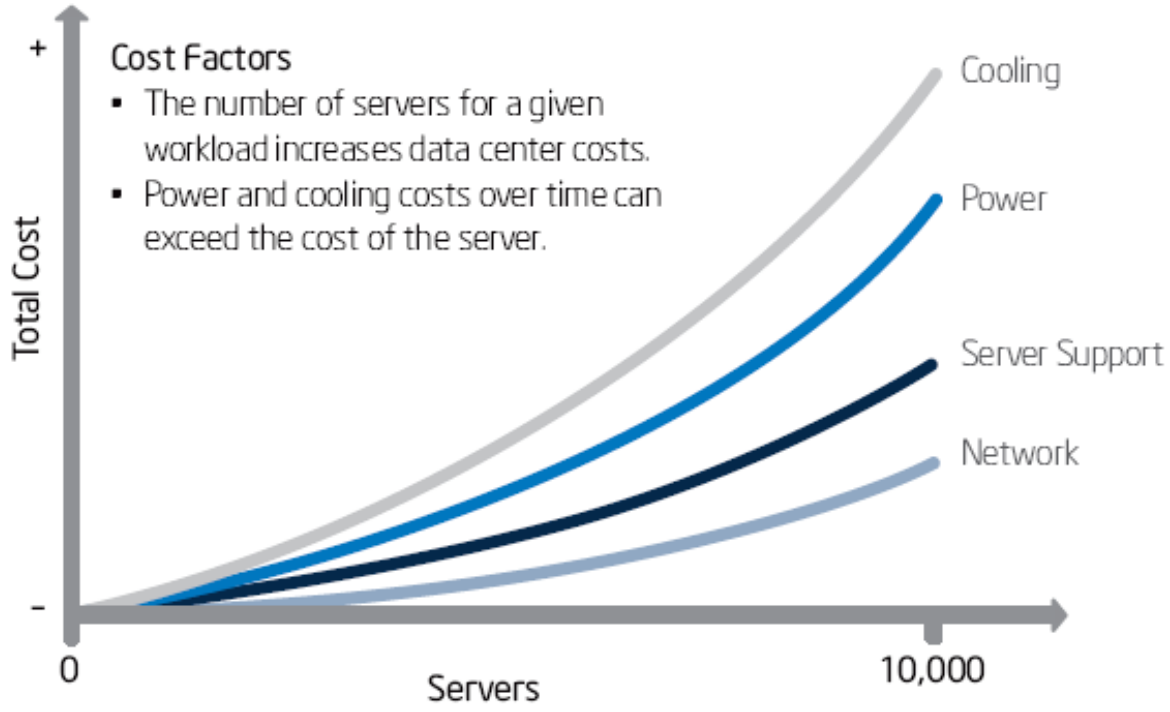

Power and cooling requirements can be the #1 and #2 costs for operating a data center. "In some cases, power costs account for 40-50% of the total data center operation budget" (Intel, 2008). It can also be a major issue beyond costs as a limiting factor for new deployments.

With the drivers of growing power costs, new workloads and systems increasing electricity requirements and a shrinking in the availability of additional power, the demand for new approaches grows. Energy consumption is predicted to remain the most dominant trend in data centers during the next five years – both from an efficiency standpoint (how to reduce consumption) and a monitoring/management stand-point (how to ensure the most efficient use of the resources at hand).

The economic downturn has caused many facilities to refurbish or expand an existing data center or to put plans for building a new one on hold. Many respondents at the Gartner conference reported that their

data centers were getting old and are not well suited to today's computing environment. In addition, growing interest in green initiatives adds to facility challenges.

Cloud Computing

Developing a public/private cloud computing strategy is another challenge that will likely be growing in the near future, as organizations are predicted by Gartner to spend more on private cloud computing services through 2012 than on public cloud services – but the decision of whether to "buy" or "build" needs to start with a deep analysis of service offerings, a strategic plan for each service and an evaluation of when external services will be available to meet enterprise needs.

Virtualization

Virtualization is one of the most significant drivers of both potential benefits and risks. Most organizations are now using some level of virtualization technology for servers and data storage in their data centers. And advanced, automated provisioning capabilities enabled in these virtualized environments, whether they are server, storage, network or desktop, bring inherent management challenges and impacts to the infrastructure. The increasing heterogeneity of technology and diversity of vendors in the data center has introduced new challenges in managing complexity, as well as offering critical components in the journey toward a real-time infrastructure and cloud computing. The proliferation of virtual servers is already an area of concern for IT teams, and there is a wider realization in the industry today that the virtual server environments must be managed in different ways from physical server environments.

Staffing

Finding and retaining talent scored relatively low in the Gartner survey, as retaining talent may not be as much of a challenge at a time when the overall unemployment rates are high and many professionals are happy to stay in their current jobs. This will likely remain the case until unemployment rates go down. Those organizations in a position to hire now are at an advantage until such time.

Solution Approaches to Data Center Resource Optimization and Cost Reduction

In order to contain costs and optimize the use of data center resources, organizations are adopting a variety of approaches -- from remote management to virtualization and storage management to high availability tools and "green IT" practices.

Remote Infrastructure Management (RIM)

The practice of managing data centers remotely, using less costly offshore resources, has grown substantially in the last three years. The development of remote monitoring tools and reduction in connectivity costs have made this possible. Today, many data center outsourcing agreements are based on the same 70% offshore / 30% onsite staffing model we have seen for applications work, resulting in a net savings in the 20% to 25% range.

Server Consolidation and Virtualization

Server and storage consolidation and virtualization can be used to improve utilization of existing hardware, thereby obviating the need to buy additional resources and minimizing floor space. According to a

recent State of the Data Center survey, 31 percent of organizations are using server virtualization and 22 percent are using storage virtualization as part of their cost-containment strategies.

Of course, because virtualization introduces complexity into the IT infrastructure as noted earlier, organizations looking to fully realize the benefits of this technology, while driving down capital costs, are advised to also implement a management framework that provides architectural flexibility and supports multiple virtualization platforms as well as physical environments.

Storage Management

A growing number of heterogeneous storage management tools automate daily and repetitive storage tasks, including RAID reconfiguration, defragmentation, file system resizing, and volume resizing. With advanced capabilities, such as centralized storage management, online configuration and administration, dynamic storage tier, dynamic multi-pathing, data migration, and local and remote replication, these solutions enable organizations to reduce both operational and capital costs across the data center. In addition, agentless storage change management tools are emerging to enable a centralized, policy-driven approach to handling storage changes and configuration drift to help reduce operational costs while requiring minimal deployment and ongoing maintenance effort.

While storage needs continue to grow, storage capacity is often underutilized. To make better use of storage resources, organizations can leverage storage management technologies. Storage Resource Management (SRM) tools, for example, enable data center staff to gain visibility into the storage environment, understand what applications are connected to each storage resource and exactly how much of the storage is actually being used by each application. Once this level of understanding is gained, organizations can make an informed decision about how to reclaim underutilized storage and also to predict future capacity requirements. 71% of respondents to the State of the Data Center survey indicated they are currently exploring SRM solutions.

Another relatively new approach focuses on reducing the duplication of data. Appropriately called Data Duplication, it can decrease the overhead associated with holding multiple copies of the same data by identifying common data and reducing copies to a single entity. This, in turn, can have a dramatic impact on the amount of disk storage required for archiving purposes as well as the number of disks required for backup purposes. Most organizations are likely to consider implementing data duplication in their efforts to maximize storage efficiency.

Other approaches to improve storage capacity utilization include thin provisioning and storage arrays, which enable capacity to be easily allocated to servers on a just-enough and just-in-time basis.

High Availability

High availability solutions such as server clustering can streamline efficiency by monitoring the operational status of the server platform components and automatically switching to a properly functioning set of components in the cluster in the event of a fault. These high availability solutions provide 99.99% or higher platform availability for critical applications.

For recovery purposes, clustering tools can be combined with replication technologies to completely automate the process of replication management and application startup without the need for complicated manual recovery procedures involving storage and application administrators. These high availability and recovery solutions also ensure increased administrator efficiency by providing a single tool for managing both physical and virtual environments.

Data Protection

Next-generation data protection tools can also be used to reduce the operational costs of protecting and archiving data as well as to meet internal SLAs and external governance requirements. With automated, unified data protection and recovery management tools that are available from a single console and work across a heterogeneous physical and virtual environment, organizations can maximize IT efficiency. A number of these tools provide for additional efficiencies through capabilities such as continuous data protection, advanced recovery of critical applications, data archiving and retention, and service-level management and compliance.

Green IT Practices

Several of the solution approaches above parallel the objectives of lowering emissions through both using less and making more efficient use of energy. Server consolidation, virtualization and storage management tools and strategies directly address these objectives.

Staffing

While many organizations respond to staffing constraints by outsourcing some IT tasks, a number of effective alternatives are also available. The most common of these strategies is to increase the automation of routine tasks. This not only reduces costs, but also frees headcount to address higher value functions. Another common strategy is to use temporary contract staff to cover workload peaks and tasks requiring periodic expertise.

About the author(s)

Dr. Pradeep K. Mukherji is President and Managing Partner at Avasant APAC and MEA

Mike Andersen is a Partner with Avasant and based at Chicago, U.S.A

About Avasant

With its global headquarters in Los Angeles, California, Avasant is a top management consulting, research, and events firm servicing global clients across the public, private, and non-profit sectors. Our talented team of consultants, lawyers and technologists average over 20 years of industry-honed experience and have conducted 1,000+ engagements in over 40 countries worldwide. Avasant drives customer value through the use of our proprietary consulting and advisory methods, which have been refined over decades of 'real-world' transaction and engagement experience. The combination of our world-class resources allows Avasant to yield superior business outcomes in three primary domains: Strategic Sourcing, Technology Optimization and Globalization Consulting.

For more details about Avasant's services and capabilities, Please visit our website at: **www.avasant.com**

Visit us for a complete list of research papers on: **www.avasant.com/research**

AVASANT

3601 N Aviation Boulevard, Suite 3000, Manhattan Beach, CA, 90266, USA
Tel: 310-643-3030
Fax: 310-643-3033
Email: contactus@avasant.com

Global Headquarters	Los Angeles
North America	Austin | New York | Chicago | San Diego | San Francisco
Asia	Beijing | Bangalore |Mumbai | Tokyo
Europe	London
Middle East & Africa	Accra

For any further communications, please contact: marketing@avasant.com