



The Art of Production Environment Engineering

By Bill Walker - SunPS ITCOPS

Sun BluePrints™ OnLine - June 2000



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 806-5810-10
Revision 01, June 2000

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun BluePrints, SunReady, SunWorld and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, SunReady, SunWorld, Sun BluePrints, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please
Recycle



Adobe PostScript

The Art of Production Environment Engineering

Abstract

Sun™ servers have entered the traditional datacenter environment, where system availability, manageability, and service availability are the key components of providing a solution to business requirements. By addressing the operating environment, as well as the application environment using an established, encompassing "production environment" lifecycle plan, robust solutions can be safely and reliably placed into production, and evolved to meet the changing and growing needs of the business.

This article will address the production environment needs of the datacenter production environment using the International Organization for Standardization (ISO) Fault, Configuration, Application, Performance, Security (FCAPS) model, the ITIL (Information Technology Infrastructure Library) framework, a basic IT reference model, and the SunReady™ Methodology's roadmap to production. This provides an encompassing view of the infrastructure requirements and timelines necessary to successfully launch, manage, and maintain a new, or refresh an existing datacenter service.

Successful production environment implementations are the product of careful planning and design. This can be considered a science as the topic of datacenter production environments has been studied and analyzed for decades. Balancing the business objectives, the explosive growth of Internet business, and the integrity and safety of the critical business services can definitely be classified as an art.

Introduction

With the exponential growth of the Internet and consumer electronic commerce on the Internet, service quality has become a key component of success. Electronic commerce and commercial portals on the Internet expose the business front office, as well as the related business back office systems to scrutiny by the direct consumer and the news media. A problem in the datacenter is no longer an “internal” problem.

Traditional “glass house” datacenter environments are generally familiar with the concepts of production environment standards for operational policies and procedures, but are unprepared for the exponential growth experienced by the fast-paced dot.com consumer markets. These glass house environments must learn to adapt to the fast-paced Internet economy, and be prepared to quickly adapt to the new model of instantaneous, direct customer contact, and an environment where exponential growth has become a standard in successful ventures. They must realize that the traditional model of multi-layered sales and distribution retail markets are being challenged by the new access model that includes direct interaction between the corporate IT systems and the customer, with competition just a customer’s mouse click away.

The dot.com startup Internet business has a related challenge. While their business model relies on exponential growth in order to survive the highly competitive Internet business environment, they often lack the datacenter policies, procedures, and experience to manage a datacenter production environment through the growth that some successful Internet businesses have displayed in recent years. By learning from the cumulative experience of their more traditional predecessors, the dot.com can implement best practices for production environment architecture, while retaining the flexibility necessary to thrive in the dot.com market.

Frameworks

An IT framework is necessary to define the scope of coverage in the production environment. The framework consists of generalized topics to be designed and implemented in the production environment. Frameworks are also useful in determining staffing and organizational requirements based on broad categories of IT services. Examples of some traditional topics include:

- Backup and Restore
- Change Control
- Security

- Performance/Capacity Management
- System Management
- Application Management
- Network Management
- Disaster Recovery
- Facilities Management

Some of the categories defined in the framework are company-wide, some are datacenter specific, some are system specific, and others are application specific. For example, the security group could manage security policy and administration for the entire company, while a facilities management resource could be designated for each of several disparate datacenters. The application management resource could be dedicated to a single server, or set of service instances throughout multiple datacenters.

The ISO FCAPS Framework

The ISO defined a framework for managed objects in a networked enterprise in their *Guidelines for the Definition of Managed Objects* contained in the *Open Systems Interconnection (OSI): Structure of Management Information* standard. The OSI structure defines five categories of managed objects, covering basic categories of enterprise support operations.

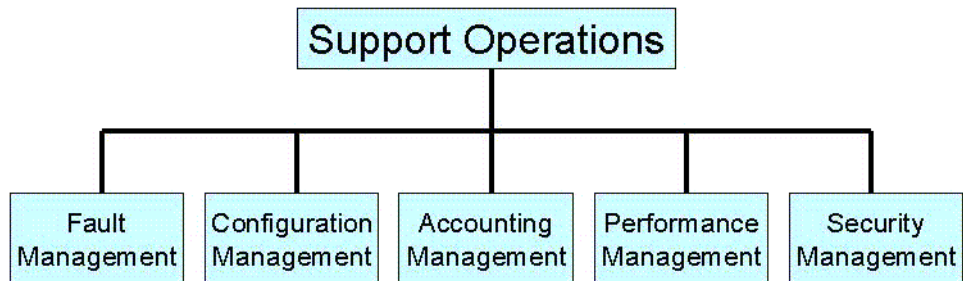


FIGURE 1 ISO FCAPS High Level IT Framework

These categories can be applied to network management systems (the intended purpose), or expanded to include the skills, tools, and capabilities in datacenter operations and management.

Extended Frameworks

Each of the categories in the IT framework can be broken down into smaller pieces, describing the responsibilities of the group providing the service associated with that category. This expanded list of IT capabilities is called the IT Extended Framework (ITEF).

Fault Management
Customer Service
Incident Management
Backup and Recover
Disaster Recovery
System Administration
Fault Protection

TABLE 1 Example IT Extended Framework

The ITEF provides an excellent checklist for datacenter capabilities, skillsets, and product coverage, as well as a framework to begin establishing the logical structure for datacenter policies and procedures, independent of operating system, server type, or application type. Several levels of criticality can be assigned, such as mission critical, business critical, business support, and incidental, in order to assign differing levels of IT support to the categories of IT services established by the ITEF.

An ITEF can be further broken down to establish functional responsibility, and provide more detailed information about the services to be provided in the ITEF. These functional responsibilities can be represented as either a table of departmental responsibilities (see TABLE 2), and explicitly declared in an IT service coverage document defining the departmental responsibilities.

Backup and Recovery
Tape Pool Management
User Data Restoration
System Backup
Data Backup Integrity Tests
Data Backup Integrity Tests
Offsite Tape Archive Rotation

TABLE 2 Example FCAPS IT Extended Framework

FCAPS Extended Framework

By expanding the FCAPS IT framework to include descriptive categories of IT service coverage, we can define an ITEF based on FCAPS. This ITEF can be represented in a tree structure, viewable as a series of webpages, with a tree structure that allows the consumer to easily navigate to the level of detail needed. This drill-down structure is key for implementing ITEF in an easily accessible, shared presentation, for inclusion in online system runbooks.

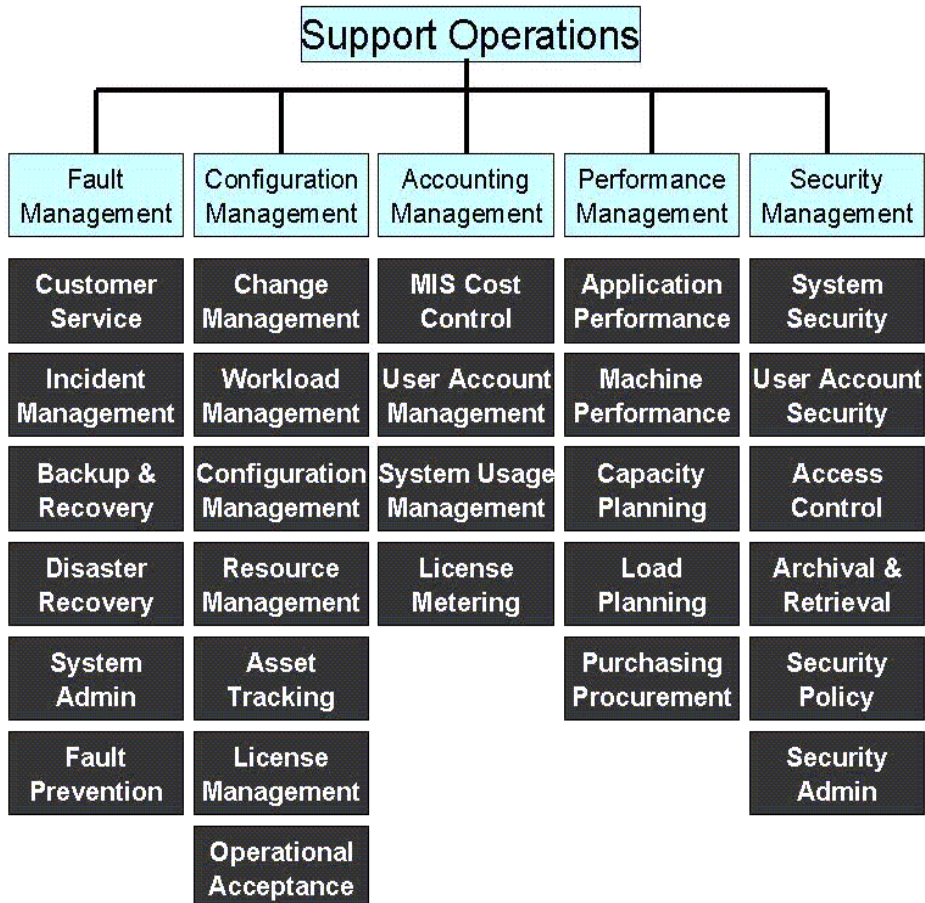


FIGURE 2 ISO FCAPS IT Extended Framework

Presentation of the logical architecture is important in presenting a unified IT management and operations strategy to both the end-user service consumer organizations, and business level management for SLA creation and service capability interpretation.

IT Service Coverage Document

The IT service coverage document should include a general statement describing the category of coverage, as well as statements describing the individual services and areas of authority in the IT environment. The service coverage document can provide assistance in tool selection, as well as staffing and training requirements. By establishing categories of coverage, the services provided by the IT department are well defined, allowing enforcement and auditing policies to be established across the IT operating environment.

This example Service Coverage Document for Change Management can serve as a guideline, but should be adjusted based on local policies and conventions:

Change Management

The change management system is the vehicle for recording, authorizing, monitoring and reporting of all changes to the Open Systems environment. The primary services and responsibilities provided are:

- Single auditable mechanism for handling all changes integrated with the enterprise change management system
- Review process for authorizing changes including all interested parties
- Process for ensuring all changes are tested and validated with tested back out procedures
- Change process reporting procedure incorporated within the SLA reporting process
- Mechanism for reviewing changes not authorized through the change management procedure.
- Control of software development and maintenance for applications and packages, including source code, executable code and parameter files, plus the provision of release management and audit control
- *Software distribution of applications, operating systems and packages to the distributed platforms in a controlled and audited manner.*

The IT Service Model

The IT Service Model (ITSM) introduces levels of criticality, levels of service, and “Big Rules” to the ITEF. This multi-level approach extends the natural “drill-down” structure for IT services.

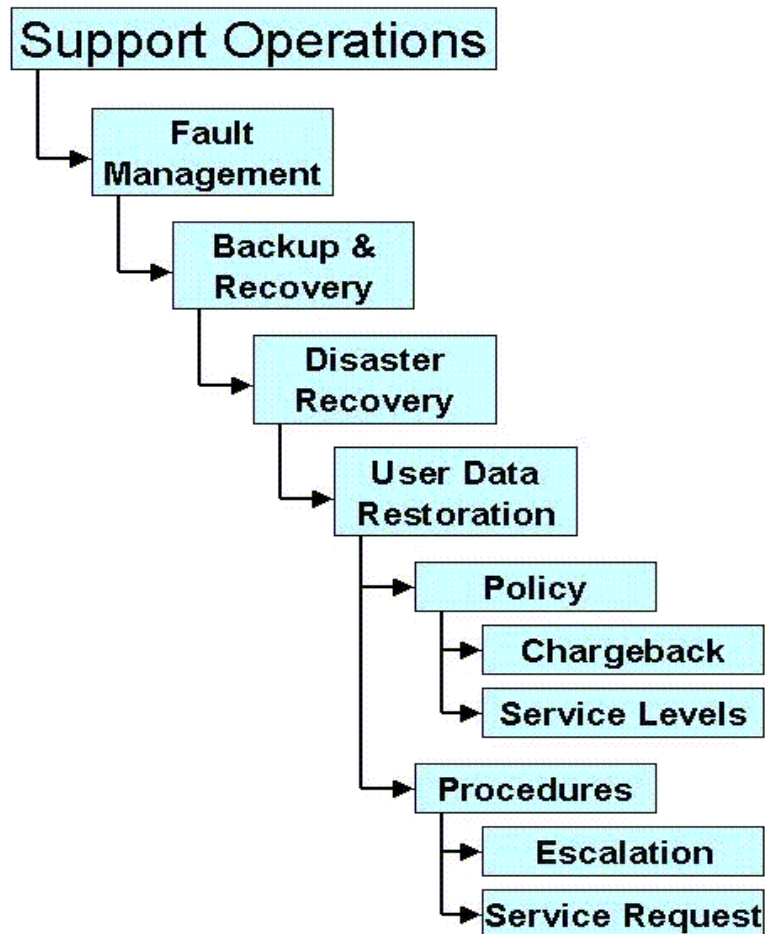


FIGURE 3 Drill-down IT Document Hierarchy Example

The ITSM can be implemented across platforms and used to create Service Level descriptions for use in SLAs between the end-user consumer organizations, and the IT resource provider organizations. Differing levels of service can be used to

produce a chargeback structure for IT services (where appropriate), as well as an automated, standardized, repeatable, and simplified SLA creation and auditing process.

This sample Service Level Description for User Data Backup can be used as a guideline, but should be modified to adhere to local datacenter policies and conventions:

Service:

User Data Backup

Service Schedule:

Maintenance Window Task

Service Levels:

Mission Critical

User data will be backed up to the central backup management system, using a full backup on a daily schedule.

Chargeback:

\$1000 + \$80 / 100 GB per month

Business Critical

User data will be backed up to the central backup management system, with weekly full backups, and daily incremental backups.

Chargeback:

\$500 + \$60 / 100 GB per month

Business Support

User data will be backed up to a locally attached tape system, with weekly full backups, and daily incremental backups

Chargeback:

\$200 + \$40 / 100 GB per month

Applying ITIL

The Central Computing and Telecommunication Agency (CCTA) of the UK began publishing volumes of their established best practices for IT operations in 1978. The 24 volumes, titled “Information Technology Infrastructure Library” (ITIL), describe the CCTA’s operational procedures, and discuss their application and logistics in a production datacenter environment.

ITIL has evolved over the years, being studied, adopted, and adapted by many IT organizations and vendors of management tools. The Examination Institute of the Netherlands (EXIN), has established certification programs granting Foundation, Practitioner, and Master levels of expertise to the categories of Service Support, Service Delivery, and IT Management.

One key benefit of ITIL is that it has helped to identify the inter-dependencies inherent in IT management organizations. For example, backups must be scheduled through the batch job scheduling system, and report errors and status through the enterprise monitoring and management system, escalating through the service desk, in order to notify the technology owner, or system administrator.

With a properly organized IT framework and extended framework, these inter-dependencies and process flows can be documented, and the procedures for implementation and operation can be integrated into the documentation describing the service.

ITIL establishes policies, and not implementations. It remains operating system, hardware platform, and application independent, but can be used as a guideline for creating the “big rules” to be used in the ITEF, which can then be implemented based on service levels, operating system, hardware, and application capabilities.

The Sun BluePrints™ program, the collective experience of the IT staff, technical publications and conferences, and the Sun Education Services can be used to supplement the existing knowledge for implementation of the established policies in the ITEF.

SunReady Program Approach

Sun Professional Services has developed a service and project methodology for applying the FCAPS framework, a framework generated from ITIL, creating a unified extended framework for IT operations. This SunReady Methodology utilizes ITIL as a guideline for best practices in mission critical environments, along with the accumulated best practices of our collective experiences. We then created sets of

“big rules” for implementation, weighting the criticality of each service in delivering the service levels promised by the customer IT department to the end-users of the specified services.

By creating this multi-level approach to each IT service, we are able to establish a scoring model, predicting the ability of the current IT infrastructure to satisfy the requirements of the end-user service level objectives. This IT capability gap analysis can then be used to create a high level architecture for the IT infrastructure necessary to supplement the current established IT operations services, before the end-user service is placed into production. The high level architecture is optimally created before the pilot phase of the application development lifecycle.

The implementation of supplemental service capabilities shadows the application development lifecycle from the architecture through the pilot, development, integration, deployment, and sustaining phases. The IT service capabilities can then be tested and assessed simultaneously with the application’s functional, stress, crash, and disaster recovery test cycles, and then promoted into the live production environment in parallel with the new end-user environment.

Attaching the IT infrastructure development lifecycle to the application development lifecycle identifies the dependencies between organizations and resources, and documents the information flow necessary to establish unified timelines for tasks and deliverable for project management. This assures that the underlying IT services required by the end-user service level objectives are in place before the end-users are introduced into the production environment.

Monitoring and Management

Service Level Management is the key to maintaining the production environment. By establishing the rules of engagement, the capabilities necessary to support the service level objectives, and the tools for use in the production environment, the entire production infrastructure can be monitored and measured. Violations of the

scope defined in the Service Level Agreements (SLA's) and measured exceptions to the agreed upon performance and management criteria can be addressed and resolved quickly, minimizing the impact on the business.

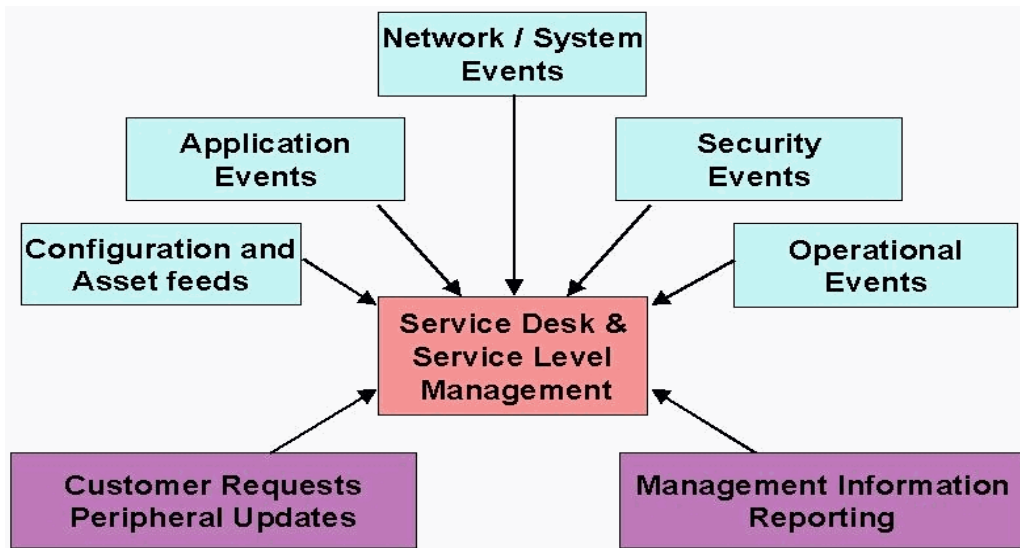


FIGURE 4 IT Business Reference Model

In this business reference model, the flow of information has been defined, with the Service Desk established as the center of the management information. The Service Desk controls information, brokers requests, audits request completion, and provides service level reporting services to the IT service management groups, as well as the business service end-user community.

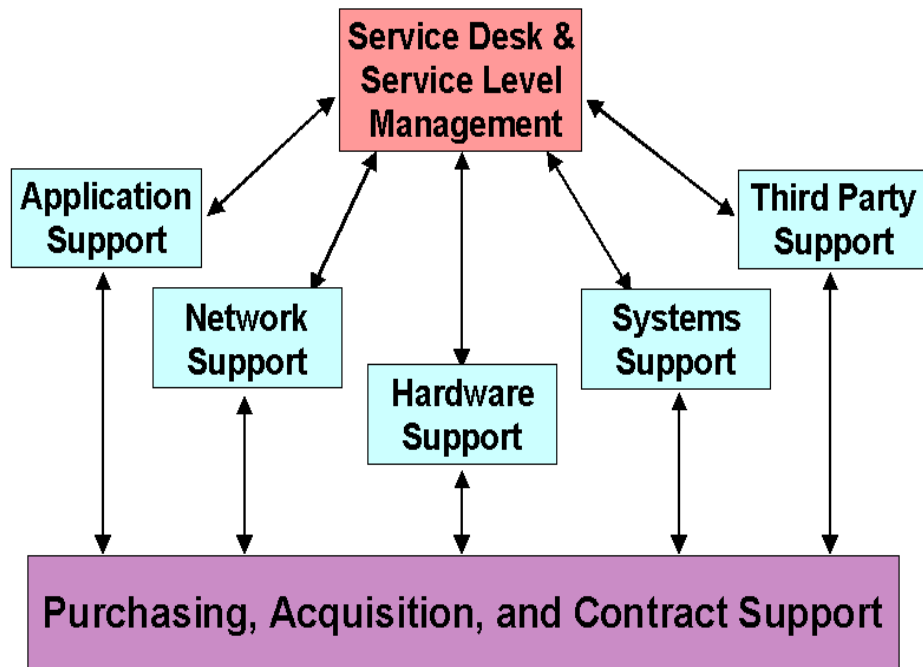


FIGURE 5 IT Business Reference Model

The Service Desk also brokers action items for the support organizations, both internal and external, and audits the completion, impact, and duration of service events. By controlling the dispatch of support activities, the affected end-user service levels can be monitored, resulting in vendor and end-user SLA and Service Level Objective (SLO) adjustments, penalty assessment, enforcement, or arbitration where necessary.

Conclusion

Production Environment Engineering may appear to be an obscure art. By studying the collective experiences of the traditional datacenter community, and applying proven best practices through an organized and well defined methodology of service coverage, it can be converted into a well understood science, with predictable quality of results. Creating reusable templates, documents and processes for design,

implementation, and integration of datacenter services in production environments can result in higher qualities of services, lower service delivery costs, and less replication of effort. Defining the framework, extended framework, reference models, and timelines, adhering to them, and enforcing them within IT production environments can help to deliver end-user services with repeatable quality of service. These define the actions necessary to keep the modern open system datacenter as well organized and well managed as the best of our traditional mainframe predecessors.

References

“The FCAPS Approach”, Geoff Payne, et al., Sun Professional Services EMEA

Sun BluePrints™ program, <http://www.sun.com/blueprints/>

“Architecting High Availability Solutions”, Wong, SunWorld™ Magazine, November 1998

Open Systems Interconnection | Structure of Management Information | Part 4: Guidelines for the Definition of Managed Objects. International Standard ISO/IEC 10165-4 : 1992

IT Infrastructure Library Online, <http://www.ccta.gov.uk/itil/>

Pink Elephant, Inc., ITIL Certification,
http://www.pdatrain.com.sg/itil/what_is_itil.htm

SunReady Methodology, <http://www.sun.com/service/sunps/sunready/>

Acknowledgements

SunPS talent and the SunReady Team were instrumental in helping me to put my thoughts into words, my express thanks to Mike Prior, Geoff Payne, Richard Croucher, Keith Hudgell, Jean-Marc Bernard, Brian Carter, Jefre Futch, and Jeffrey Lucove, along with the cast of thousands behind the scenes who all contribute to Sun’s success.

Author's Bio: Bill Walker

As an IT professional in the ITCOPS organization of Sun Microsystems' Professional Services Division, Mr. Walker has served as an IT Architect, Senior Project Engineer, and Project Technical Manager for strategic clients, both domestic and international. Mr. Walker has assumed key roles in projects including Capacity and Performance Planning, Server Consolidations, Messaging System Analysis and Conversion, IT Architecture Assessment, and System Management Software Implementation, concentrating on the Sun Enterprise E10000 and High End Data Center environments. In depth technical expertise, combined with a real world view of business implementations provide project lifecycle experience necessary for business implementation engineering in the high-end data center environment.