

Sun Datacenter Express

Solving The Customer's Red Shift Problem

**White Paper
November 2007**

Table of Contents

Introduction Red Shift: A Brief History.....	3
Hyper-growth Market Segments (Red-shifting).....	3
ΣBW (Sum-of-Bandwidth).....	3
High Performance Computing (HPC).....	3
*prise (or "Star-prise").....	3
Traditional Computing Markets (Blue-shifting).....	3
Consequences & Industry Commentary.....	4
Service Summary.....	4
Service Architecture.....	6
Standard Architectures.....	7
HPC Compute Module.....	8
HPC Scheduler Module.....	9
HPC Compiler Module.....	9
HPC Licensing Module.....	9
Service Library.....	10
SunSpectrum Support.....	10
Sun Preventive Services.....	10
Sun Managed Services	10
Sun Managed Operations.....	11
Sun Interim Operations Management	11
Sun i-Runbook Service	11
Customer Training.....	11
Customized, Factory Assembled and Tested.....	12
Conclusion.....	12

Introduction Red Shift: A Brief History

Chapter 1

Red shift is a techno-economic theory suggesting hyper-segmentation of IT markets based on whether individual computing needs are over or under-served by Moore's Law, which predicts the doubling of computing transistors (and therefore roughly computing power) every two years.

The theory, proposed and named by Sun Microsystems CTO Greg Papadopoulos, categorizes a series of high growth markets (Red-shifting) while predicting slower GDP-driven growth in traditional computing markets (Blue-shifting).¹

Hyper-growth Market Segments (Red-shifting)

According to the Red Shift theory, "Red-shifting" applications are those applications whose requirements for compute power exceeds Moore's Law, in these markets, customers are running out of data-centre real estate, power and cooling infrastructure.

Red Shift customers can be identified by the following characteristics:

ΣBW (Sum-of-Bandwidth)

Companies who are driving heavy Internet traffic. This includes popular web portals like Google, Yahoo and MSN. It also includes telecommunications, television over IP and online games.

High Performance Computing (HPC)

Companies doing complex simulations such as weather, stock market or drug design simulations. This is a generally elastic market because businesses frequently spend every "available" dollar budgeted for IT.

*prise (or "Star-prise")

Companies aggregating traditional computing applications and offering them as services, typically in the form of Software as a Service (SaaS). For example, companies deploying CRM solutions are over-served by Moore's Law, but for those companies that aggregate their applications and offer them as a service, such as salesforce.com, are seeing demand growing faster than Moore's Law.

Traditional Computing Markets (Blue-shifting)

Red Shift theory suggests that traditional computing markets, such as those serving Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM) applications, have reached relative saturation in industrialized nations. Thereafter, proponents argue further market growth will closely follow Gross Domestic Product (GDP) growth, which typically remains under 10% for most countries annually. Given that Moore's Law continues to predict accurately the rate of computing transistor growth, which roughly translates into computing power doubling every two years, the Red-shift theory suggests that traditional computing markets will ultimately contract as a percentage of computing expenditures over time.

Functionally, this means "Blue-shifting" customers can satisfy computing requirement growth by swapping in faster processors without increasing the absolute number of computing systems.

¹ Wikipedia.org

Consequences & Industry Commentary

Greg Papadopoulos argues that while traditional computing markets remain the dominant source of revenue through the late 2000s, a shift to hyper-growth markets will inevitably occur. When that shift occurs, he argues computing (but not computers) will be a utility, and differentiation in the IT market will be based upon a company's ability to deliver computing at massive scale, efficiently and with predictable service levels, much like electricity today.

According to International Data Corporation (IDC) vice president Matthew Eastwood, "IDC believes that the IT market is in a period of hyper segmentation. There is a class of customer that is Moore's law driven and as price performance gains continue, IDC believes that these organizations will accelerate their consumption of IT infrastructure"².

Service Summary

Chapter 2

In order to service the requirements of the Red Shift customer Sun Microsystems needs to be able to deploy greater volumes of horizontally scalable architecture in faster time scales. Thus the concept of the Datacenter Express Service (DES) was born. In addition to rapidly provisioning production applications on the new systems, the service considers what customers need in terms of post-deployment management services. With Datacenter Express, Sun is helping organizations move towards managing the essential services that drive the business by providing robust, flexible and manageable infrastructures.

Datacenter Express provides the basic building blocks to enable Sun to:

- Reduce initial installation time scales and implementation issues by up to 80%
- Accommodate large-volume, highly repeatable deployments
- Create opportunities to transition purchasing decisions from CapEx to OpEx
- Reduce configuration complexity and provide project and life cycle management services
- Simplify multiple vendor management and resource coordination
- Free up internal IT resources to focus on core business needs and new services.

Datacenter Express is an end to end offering for customers who require a single service to design, rapidly deploy and manage IT infrastructures. The service uses Sun Microsystems internal Customer Ready Systems (CRS) to install, load and pre-configure DES infrastructure at a CRS facility then ship the systems to the customers site. This enables rapid deployment of consistently configured horizontally scalable architectures with minimal disruption to the customers daily operations.

The Datacenter Express Service is based on an architecture and implementation methodology that has been optimized for the rapid deployment of standard building blocks for the infrastructure, operations management, and service management elements. In addition, it leverages many of Sun's existing services and expertise such as Spectrum Support, Sun Preventive Services, Interim Operations Management, Managed Operations, Operational Definition Workshop and Professional Services Consulting. Sun can design custom solutions to meet unique customer business requirements and IT needs which will leverage standard configurations and services.

This service transforms the way Sun delivers infrastructure by enabling customers to select the most appropriate

² IDC: Dell Reacts to Critical Market Shift and Announces Formation of New Data Center Solutions Division
March 28, 2007 [IDC link](#)

Datacenter Express Service option. The three options are listed below in detail but each option allows the customer a choice of support elements; designed to progressively ease the burden of installing and managing large infrastructures.

The Datacenter Express service utilizes existing Sun services and provides a flexible integrated offering to meet customers specific requirements. The service innovation has been realized in the way Sun packages these services making it easier to respond to customers requirements for rapid delivery and simple management of large complex configurations and is based on Sun's experience of delivering many such IT architectures.

The following table shows the range and scope of existing services that are included in the Datacenter Express Service:

Availability	Service Name	Option 1	Option 2	Option 3
Required	DES Architecture Workshop	x	x	x
Required	DES Operational Definition Workshop			x
Recommended	DES Site Survey	x	x	x
Required	DES Implementation Service	x	x	x
Required	Customer Ready Systems	x	x	x
Required	SunSpectrum support	x	x	x
Recommended	Sun Preventive Services		x	x
Recommended	Sun Managed Operations (on-site, remote or combination)			x
Recommended	Sun Interim Operations Management (IOM)			x
Optional	Sun i-Runbook Service	x	x	x
Recommended	Service Delivery Management			x
Optional	Remanufactured Equipment	x	x	x
Optional	Financial Services Option (leasing)	x	x	x

Table 1: Sun Datacenter Express Service Options

Service Architecture

Chapter 3

The Datacenter Express service is based on Sun's tried and tested Architect, Implement and Manage (AIM) Methodology. The three phases of any Datacenter Express deployment have clearly defined deliverables and handover points.

The **Architecture phase** consists of one or more workshops to gather the information required to design a solution to meet customer specific requirements. During the Architecture phase customer business or technical requirements and required service levels are ascertained and documented. These requirements are converted into an architecture design, a build specification and a bill of materials as well as an implementation schedule. At the closure of this phase the customer is presented with a proposal for the systems, the implementation plan, management services and the architecture design.

The **Implementation phase** takes place at two locations, the CRS factory and at the customer's site. CRS will unpack, build, rack and configure the equipment according to the architecture design. Sun and 3rd party equipment can be combined to meet customer specific requirements. CRS build and test the equipment in their facilities before shipping the racked systems to the customer site. Once on site, Sun deploys the systems as agreed during the Architecture phase. On completion of the Implementation phase, Sun will hand the systems over to the systems management team.

In the **Manage phase** the focus is the ongoing support, operation and management of the systems. The output of the Operational Definition Workshop is central, in defining the key service requirements. These in turn reflect the levels of service required, and the options of in-house or external service delivery. This needs to integrate with exiting management infrastructure, and to take account of current skills. Appropriate selections from Sun's range of Support, Managed and Educational services will provide the additional capabilities that may be needed.

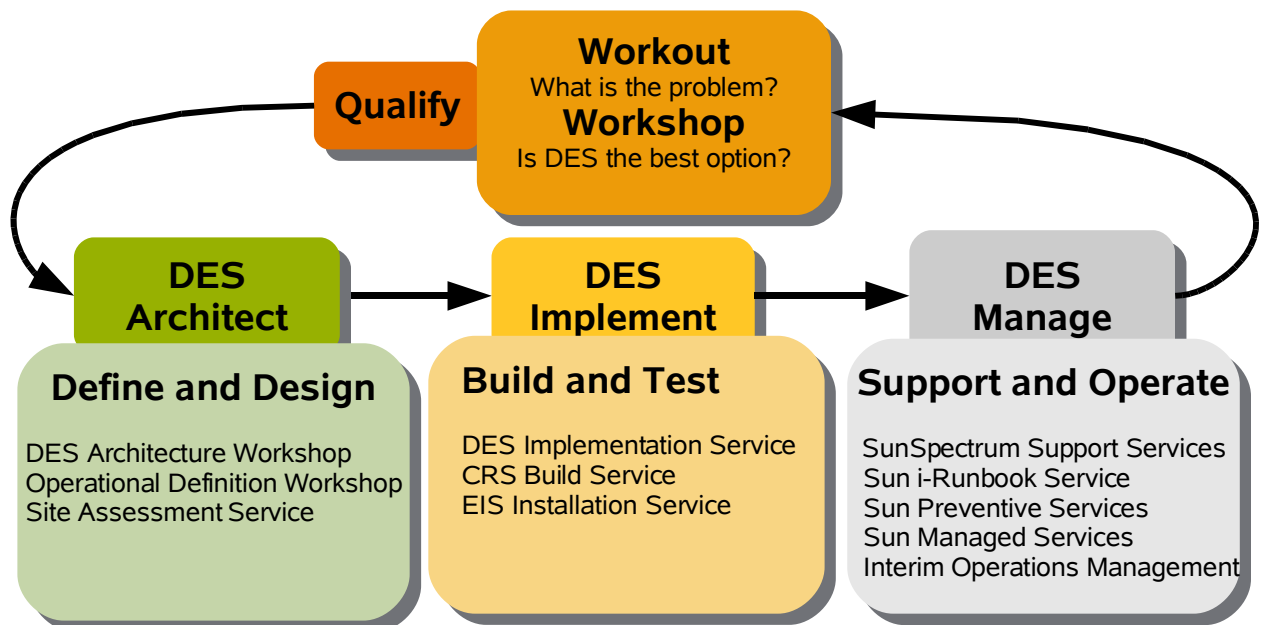


Figure 1: Datacenter Express Services process overview

Standard Architectures

Chapter 4

Sun has developed a process to decompose architectures into standard modules each designed to solve specific business problems. At its launch the Sun Datacenter Express Service will offer solutions in three distinct areas HPC (High Performance Computing), Storage and Web. These offerings will provide customers with clearly defined solution components which are designed to reduce deployment complexity and lower risk using the Datacenter Express Service.

These offerings are based on “Preferred Platforms” and utilise a concept developed by CRS called Service Delivery Modules. These modules are based on CRS experience over many years of building large and complex systems for customers. The modular approach allows CRS to quickly build complex solutions from standard components or modules.

Service Delivery Modules

A Service Delivery Module (SDM) clearly defines the scope of a delivery service. It is the implementation of a technology which meets a customer requirement to provide the expected functionality.

Each of the Service Delivery Modules consist of:

- Technology/Scope Definition Document
- Customer Requirements Questionnaire
- List of Deliverables
- Implementation Plan
- Execution
- As Built Documentation

The benefits of this approach are numerous both to the customer and to Sun. Each Service Delivery Module has a clearly defined scope and provides a single IT service. SDMs can be combined to provide a range of services. SDMs may include hardware, software or combinations of both. A clear list of deliverables identifies the scope of the SDM and opportunities for re-use of the SDM. In addition this integrated modular approach simplifies the architecture model for complex solution engagements. The Datacenter Express service benefits the customer because the required solution can be rapidly built and deployed.

Many of the Service Delivery Modules are generic and may be used to support different customer requirements. However, for some solutions Service Delivery Modules will be designed specifically to augment the platform definition. The platform definition defines the basic hardware, firmware and software configuration required in order to make the stand-alone platform ready for the implementation of solution specific Service Delivery Modules.

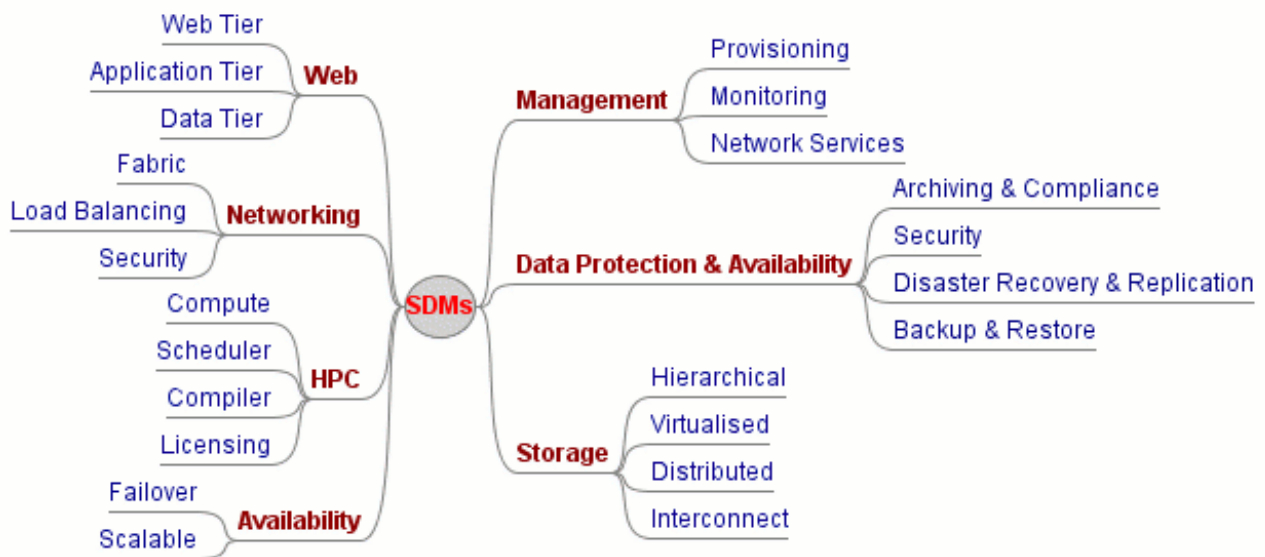


Figure 2: Datacenter Express Service Delivery Modules

Generic Service Delivery modules covering Management, Networking and Storage along with solution specific modules are used to complete the Solution Definition. Generic SDMs include:

- Platform Definition Module - defines the basic hardware, firmware and software configuration
- Provisioning Module – covers OS deployment and post installation configuration
- Monitoring Module – covers hardware and service monitoring
- Network Fabric Module – covers the configuration of all networks
- Distributed Storage Module – covers the configuration of storage devices configured with distributed or parallel file systems.

As an example of the SDMs in action we will work through the HPC solution. The High Performance Computing modules relate to all aspects of job submission, control and execution and extends to capabilities such as the compiler and licensing servers. In addition to the Platform Definition Module there are a number of HPC specific modules, these are:

- HPC Compute Module
- HPC Scheduler Module
- HPC Compiler Module
- HPC Licensing Module

HPC Compute Module

This HPC specific module provides the compute framework for the Grid including the installation and configuration of software and protocols which are required for job execution.

In scope:

- Message Passing Interface (MPI) software stack
- Infiniband Vendor supplied MPI

- OpenMPI for Gigabit Ethernet

Out of scope:

- Compute Node provisioning. Covered by Provisioning SDM
- Hardware Drivers. Covered by Platform Definition SDM
- 'Business Ready' configuration, including deployment of customer's application.

HPC Scheduler Module

The HPC Scheduler module provides the ability to submit jobs, monitor their execution and administer the batch/queue system used for scheduling jobs to the Grid.

In scope:

- Job Scheduling environment
- Job Submission and Administration nodes
- Client (Compute Node) daemons and services

Out of scope:

- Scheduling Node provisioning. Covered by Provisioning SDM
- Hardware Monitoring and Control. Covered by Monitoring SDM

For the Datacenter Express Service the initial offering will be limited to:

- Sun N1 Grid Engine

HPC Compiler Module

The HPC compiler module provides all of the necessary tools and libraries required to compile code for execution. Also includes availability of compiler runtime/libraries across the solution.

In scope:

- Compiler environment
- Libraries

Out of scope:

- Compiler Node provisioning. Covered by Provisioning SDM
- Job Submission and Administration. Covered by the HPC Scheduler Module

For the Datacenter Express Service the initial offering will be limited to:

- Sun Studio
- GNU gcc

HPC Licensing Module

The HPC Licensing module provides the necessary service for compliance with customer's application license requirements.

In scope:

- License Server environment

Out of scope:

- License Node provisioning. Covered by Provisioning SDM
- Job Submission and Administration. Covered by HPC Scheduling SDM
- Compiler and Libraries. Covered by the HPC Compiler SDM

Service Library

Chapter 5

In addition to the CRS build and configuration service via the Service Delivery Modules the Datacenter Express Service calls on existing elements from the Services Library. The library consists of services from Sun's Managed, Professional, Learning and Support Services organizations. Sun's established range of services can be combined in a logical manner to provide the correct solution to a customer's business problem.

The services library consists of:

- SunSpectrum Support
- Sun Preventive Services
- Sun Managed Services
 - Sun Managed Operations
 - Sun Interim Operations Management
- Sun i-Runbook Service
- Customer Training

SunSpectrum Support

When customers choose SunSpectrum Support, customers get more than the typical service relationship. With one overarching support agreement, Sun provides customers with virtually everything customers need: integrated hardware and operating system support for customers systems; hardware and firmware support for customers data storage; operating system-specific coverage; and options for production system software.

Customers maintain control by selecting the precise level of service customers need for individual system performance. As customers move up to higher level plans, coverage hours broaden, response time commitments quicken, and issue escalation is accelerated. Customers choose any or all of the Sun Service Plans they need, and relax. Support coverage is all under one Sun.

Sun Preventive Services

Sun Preventive Services provides a realistic assessment of customer's data center's health based on measurable risk indicators, plus a treatment plan and incentives for maintaining improvements. Sun's multidimensional, preventive model changes data center support from "hoping nothing goes wrong" to planning and management that can result in operational excellence and give customers a competitive advantage.

Sun Managed Services

The Sun Managed Services portfolio of offerings allows Sun Microsystems to manage customer-selected aspects of day-to-day IT operations.

Unlike many of our competitors, Sun's approach is to offer customers a selective-sourcing delivery model. Here, in addition to having chosen aspects of their day-to-day IT operations, customers have the flexibility to choose the service features they need and for the duration that suits them, all while leveraging Sun's expertise and technology. Sun's services are based on an ITIL-based framework, reassuring the customer that Sun is process-oriented and driven to do things right the first time.

Combined with Professional Services engagements and Support Services like SunSpectrum, Sun Microsystems can provide customers with a holistic solution to operate and manage IT infrastructures in a very focused way that allows the Customer to maintain choice, control, and flexibility.

Sun Managed Operations

Delivers long-range IT operations management that allows the customer to align their IT competencies with their critical business objectives. Sun Managed Ops offers remote monitoring and management for all components of a Customer's IT infrastructure, including systems at the OS level, third-party, custom, and SunRay applications, databases, networks, and the web. The service creates customized solutions for a Customer environment based on the components they've selected for Sun to monitor and/or manage.

Sun Interim Operations Management

A customized service providing on-site management of a customer's IT operations. Sun technical experts and operational consultants work to help reduce operational risk during transitional periods such as major environment upgrades or technology refreshes. During the engagement, a highly qualified Sun team improves and manages the day-to-day operation of the Sun platform. When the engagement ends, Sun transitions data center operation to whatever long term data center management options the Customer prefers.

Sun i-Runbook Service

The Sun i-Runbook service creates a knowledge power-base for customers organization, allowing an IT team to find the vital information they need to manage the data center using the latest Sun preferred practice advice from a central source – accessed within seconds via a simple web interface.

The service enables customers operations personnel, system administrators and IT managers to better understand and manage customers Sun and third-party hardware and software. This can help lead to a reduction in errors, faster problem-resolution, widespread adoption of preferred practices and grow technical knowledge among the team.

Customer Training

Sun offers a wide range of IT services from consulting to courseware to certification, to improve expertise and strengthen performance of an IT team.

Customized, Factory Assembled and Tested

Chapter 6

The Sun Datacenter Express Service enables faster, more efficient, higher quality build and testing of the customer's configuration—without disruption to their datacenter by utilizing Sun's Customer Ready Services facility.

CRS provides the interim hosting facilities for all build, test and operations management activities the solution is then packed and shipped to the customer datacenter where it is installed, verified and integrated into the customer data center. Typically from receipt of approved architecture to delivery at a customers site takes as little as 4 weeks.

A recent Sun CRS program engagement achieved dramatic results in all aspects of system deployment.

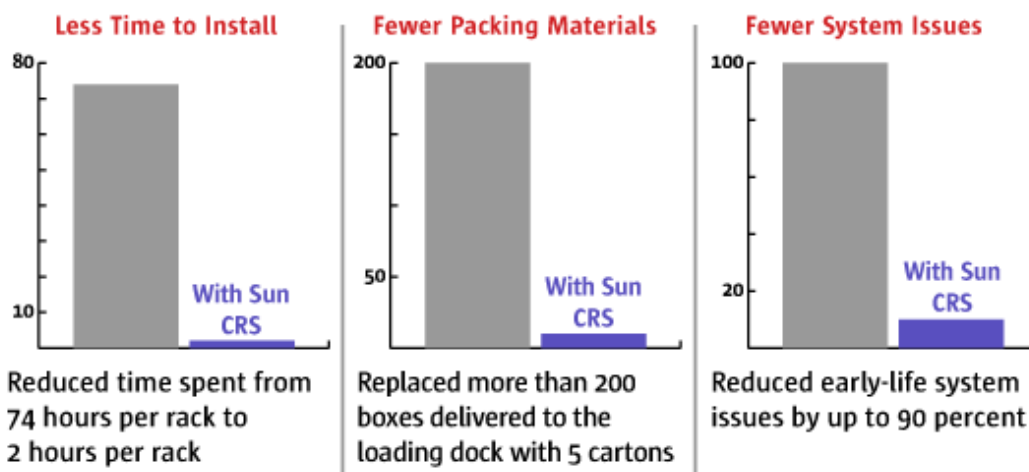


Figure 3: Service Deployment Time and Reliability Improvements

Conclusion

Chapter 7

In conclusion Red Shift companies who are driving heavy Internet traffic such as the popular web portals Google, Yahoo and MSN or companies doing complex simulations such as weather, stock market or drug design simulations are perfect candidates for this service.

Sun's new Datacenter Express Service provides rapid, high quality deployment of complex customer solutions to service the ever changing requirements of the Red Shift customer. With the Datacenter Express Service, Sun is helping organizations move towards managing the essential services that drive the business by providing robust, flexible and manageable infrastructure. Sun's new Datacenter Express Service reduces initial installation time scales and implementation issues by up to 80%. It is the perfect service to support the Red Shift customer base but it is important not to forget that any customer who values reduced time to deployment and high quality services at a premium price would also benefit from Sun's Datacenter Express Service.

Sun Microsystems, Inc. Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com

©2007 Sun Microsystems, Inc. All rights reserved Sun, Sun Microsystems, the Sun logo, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

SunWin #519247