

Ideal Platforms for Optimizing IT Workloads

TRIED AND TESTED SOLUTIONS TO ACCELERATE IT AND BUSINESS TRANSFORMATION

First Edition

Marty Poniatowski

HPE Press
660 4th Street, #802
San Francisco, CA 94107

Ideal Platforms for Optimizing IT Workloads

Tried and tested solutions to accelerate IT and business transformation

Marty Poniatoski

© 2017 Hewlett Packard Enterprise Development LP.

Published by:

Hewlett Packard Enterprise Press
660 4th Street, #802
San Francisco, CA 94107

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

ISBN: 978-1-942741-54-1

WARNING AND DISCLAIMER

This book provides information about optimizing IT workloads using various technologies and solutions. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The author and Hewlett Packard Enterprise Press shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

The opinions expressed in this book belong to the author and are not necessarily those of Hewlett Packard Enterprise Development LP.

GOVERNMENT AND EDUCATION SALES

This publisher offers discounts on this book when ordered in quantity for bulk purchases, which may include electronic versions. For more information, please contact U.S. Government and Education Sales 1-855-447-2665 or email sales@hpepressbooks.com.

Feedback Information

At HPE Press, our goal is to create in-depth reference books of the best quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the expertise of members from the professional technical community.

Readers' feedback is a continuation of the process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at hpepress@epac.com. Please make sure to include the book title and ISBN in your message.

We appreciate your feedback.

Publisher: Hewlett Packard Enterprise Press

HPE Press Program Manager: Michael Bishop

CONTENTS

| | |
|--|-------------|
| Foreword | xiii |
| Introduction | xv |
| 1 Anatomy of Workloads | 1 |
| The evolution of workloads | 2 |
| Workload categories..... | 4 |
| The Internet of Things (IoT) as another workload category | 6 |
| Workloads Framework..... | 7 |
| Summary..... | 9 |
| 2 High Performance Computing: Business Intelligence | 11 |
| Solution overview..... | 12 |
| From business goals to PoC to production | 14 |
| Hardware inventory..... | 16 |
| Software inventory..... | 19 |
| Deployment..... | 19 |
| Summary..... | 20 |
| 3 High Performance Computing: Scientific Research | 21 |
| Solution overview..... | 22 |
| How did we arrive at this solution? | 23 |
| Server space | 24 |
| Server power..... | 24 |
| Server cooling | 25 |
| Hardware inventory | 25 |
| Implementation project plan..... | 28 |
| Solution validation and success criteria | 28 |
| Summary..... | 28 |
| 4 Scale-Out Workload | 29 |
| Scaling a single-threaded application | 31 |
| How did we arrive at this solution? | 32 |
| Hardware inventory | 33 |
| Software inventory | 34 |
| Implementation project plan..... | 35 |

| | |
|---|-----------|
| Expertise and skills | 36 |
| Solution validation and success criteria | 36 |
| Application architecture | 37 |
| Test application description | 37 |
| Test workload data/results | 38 |
| Analysis and recommendations | 39 |
| Deployment overview | 39 |
| Management tool overview and comparison | 40 |
| Summary | 41 |
| 5 Scalable Storage | 43 |
| Case #1: Object storage for unstructured data | 43 |
| Why object storage? | 44 |
| Solution overview | 44 |
| Hardware inventory | 46 |
| Storage server (16 Total) | 46 |
| Connector server (2 Total) | 47 |
| Supervisor server (1 Total) | 48 |
| Object storage software inventory | 48 |
| Sizing calculations | 49 |
| Key features of object storage | 49 |
| Erasure coding | 49 |
| Scale-out | 50 |
| System management | 51 |
| Deployment overview | 51 |
| Case #2—Two-site object storage | 52 |
| Active/Passive object storage | 52 |
| Solution design | 53 |
| Summary | 55 |
| 6 Archiving Large Files | 57 |
| Case #1—Archiving large files to disk | 57 |
| Solution overview | 58 |
| How did we arrive at this solution? | 59 |
| Server and storage | 59 |
| Hardware RAID settings | 61 |
| Controller selection | 61 |
| Hardware inventory | 64 |
| Implementation project plan | 66 |

| | | |
|----------|--|-----------|
| | Solution validation and success criteria | 66 |
| | Looking ahead..... | 67 |
| | Case #2—Replacing tape with object storage for file archive..... | 67 |
| | Solution overview..... | 69 |
| | Solution details | 70 |
| | Summary..... | 72 |
| 7 | Hosted Desktop Infrastructure (HDI) | 73 |
| | Solution overview..... | 73 |
| | How did we arrive at this solution? | 76 |
| | Solution inventory | 77 |
| | Client configuration..... | 78 |
| | Cartridge configuration | 78 |
| | Citrix software | 79 |
| | Moonshot software..... | 79 |
| | Moonshot power supplies | 80 |
| | System management | 80 |
| | Summary..... | 80 |
| 8 | Virtual Desktop Infrastructure (VDI) | 81 |
| | Solution overview..... | 81 |
| | How did we arrive at this solution? | 83 |
| | Processor selection | 84 |
| | StoreVirtual VSA and Hypervisor Compute Considerations | 85 |
| | StoreVirtual VSA performance and capacity | 85 |
| | Solution inventory | 86 |
| | Software for VDI environment | 88 |
| | Summary..... | 88 |
| 9 | engineering Virtual Desktop Infrastructure (eVDI)..... | 89 |
| | Why eVDI | 89 |
| | Business benefits of eVDI | 89 |
| | Solution overview..... | 90 |
| | Datacenter considerations..... | 95 |
| | How did we arrive at this solution? | 96 |
| | Processor selection | 99 |
| | Disaster Recovery (DR) considerations | 99 |
| | Solution inventory | 99 |
| | Software components | 100 |
| | Summary..... | 100 |

| | |
|--|------------|
| 10 Citrix Implementation on HPE Hyper Converged 250..... | 101 |
| Solution overview..... | 102 |
| Design for Multiple locations..... | 103 |
| How did we arrive at this solution?..... | 104 |
| HPE StoreVirtual technology..... | 105 |
| Citrix XenApp Environment..... | 105 |
| Hardware and software requirements..... | 107 |
| Hardware requirements..... | 107 |
| Management software requirements..... | 110 |
| Microsoft licensing requirements..... | 110 |
| General requirements..... | 110 |
| Summary..... | 110 |
| | |
| 11 EMR Software on HPE ConvergedSystem..... | 111 |
| ConvergedSystem 700..... | 111 |
| Solution overview..... | 113 |
| Design considerations..... | 116 |
| Compute platform design considerations..... | 116 |
| Storage design considerations..... | 117 |
| Management and monitoring..... | 120 |
| Implementation and support..... | 121 |
| Combined hardware and software inventory..... | 121 |
| Summary..... | 126 |
| | |
| 12 Hyper Converged System Running Multiple Workloads..... | 127 |
| Solution overview..... | 127 |
| Hardware inventory..... | 129 |
| Software inventory..... | 132 |
| Implementation project plan..... | 133 |
| Solution validation and success criteria..... | 133 |
| Summary..... | 136 |
| | |
| 13 SAP HANA..... | 137 |
| SAP HANA deployment for nonproduction environment..... | 137 |
| Sizing and design..... | 138 |
| Solution overview..... | 139 |
| SAP HANA infrastructure input sizing summary..... | 139 |
| SAP infrastructure design assumptions..... | 140 |
| Solution components..... | 140 |

| | |
|--|------------|
| Hardware inventory | 142 |
| Deployment overview | 144 |
| Summary..... | 144 |
| 14 Microsoft SQL Server Scale-Up Workload..... | 145 |
| Summary of Microsoft workload..... | 145 |
| Solution overview..... | 146 |
| How did we arrive at this solution? | 147 |
| Hardware inventory | 149 |
| Software inventory | 150 |
| Implementation plan..... | 151 |
| Expertise and skills | 152 |
| Solution validation and success criteria | 152 |
| Application architecture | 153 |
| Application testing | 153 |
| Analysis..... | 156 |
| Summary..... | 156 |
| 15 Backup and Recovery of Virtual Machines..... | 157 |
| Solution overview..... | 157 |
| How did we arrive at this solution? | 160 |
| Hardware inventory | 161 |
| Software inventory | 163 |
| Implementation steps | 163 |
| Fiber channel step-by-step setup instructions | 164 |
| iSCSI step-by-step setup instructions | 165 |
| Summary..... | 167 |
| 16 Hyper Converged Backup and Recovery..... | 169 |
| Solution overview..... | 169 |
| How did we arrive at this solution? | 170 |
| Hardware inventory | 171 |
| Software inventory | 175 |
| Solution validation and success criteria | 176 |
| Implementation, deployment, and testing | 176 |
| Test workload results..... | 176 |
| Summary..... | 177 |

| | |
|--|------------|
| 17 Disaster Recovery..... | 179 |
| Current and new configurations | 180 |
| The new solution overview | 181 |
| How did we arrive at this solution? | 183 |
| Hardware and software inventory | 186 |
| Solution validation and success criteria | 189 |
| Deployment overview | 190 |
| Validation | 190 |
| Summary..... | 191 |
| | |
| 18 An Asymmetric Approach to Hadoop | 193 |
| Design Factors | 193 |
| Solution overview..... | 195 |
| How did we arrive at this solution? | 197 |
| Compute platform selection..... | 198 |
| Compute node design | 199 |
| Drive selection | 199 |
| Network setup | 199 |
| Management tool selection | 200 |
| Management node design | 200 |
| Proactive care advanced selection | 200 |
| Hardware inventory | 201 |
| Implementation project plan..... | 206 |
| Summary..... | 206 |
| | |
| 19 Hybrid IT: The “Right Mix” | 209 |
| Solution overview..... | 210 |
| Solution components | 211 |
| Hyper Converged 380 | 211 |
| Helion CloudSystem 10 Enterprise..... | 212 |
| Cloud Service Automation | 213 |
| Operations Orchestration | 214 |
| Helion OpenStack..... | 215 |
| Helion Stackato | 216 |
| Solution inventory | 217 |
| Services and education..... | 220 |
| Summary..... | 221 |

| | |
|---|------------|
| 20 The Power of Connection: The Internet of Things | 223 |
| The opportunity with connected, intelligent things..... | 224 |
| Common IoT architecture elements..... | 226 |
| Device connection and data acquisition (Stage 1) | 227 |
| Data aggregation (Stage 2) | 228 |
| Edge analysis (Stage 3) | 229 |
| Deep integration with IT systems and other data sources (Stage 4) | 231 |
| Security considerations (supporting all stages) | 231 |
| The potential of a connected future | 232 |
| 21 Memory-Driven Compute and Composable Infrastructure | 235 |
| Memory-driven computing..... | 235 |
| HPE Synergy..... | 239 |
| Management of Synergy | 242 |
| Epilogue | 243 |
| Index..... | 245 |

Foreword

The industries and customers we serve are going through a significant digital transformation in both their business models and the markets they participate in. These transformations are bringing new opportunities, and are enabling our customers to drive the right outcomes for their business.

Technology has progressed to a point where you can target your environment to meet specific workload needs - and these workloads vary widely! You may have, for instance, a complex problem that will scale-out to thousands of nodes and requires constant tuning, or at the other end of the spectrum, a remote office that needs minimal resources and there is no one to maintain the system.

In this book, HPE experts cover one of the most important technical topics that exists today: Optimizing your environment for individual workloads. The book also provides a rare opportunity to review several tried and tested deployments that have produced excellent results, all based on a mature set of solutions driven by experts in the field with real examples and designs.

From a broader perspective, this book is not simply about workloads. It also demonstrates the fact that Hewlett Packard Enterprise offers not just one or two proprietary platforms but a whole suite of open and flexible IT solutions with innovations such as HPE Synergy and Aruba ready to bridge the evolution to Composable Infrastructure, multi-cloud environments, and the intelligent edge. This aligns with HPE's vision that, in the future, the world will be hybrid, everything will be software-defined, and the edge will explode. The examples you will find in this book showcase the HPE strategy – to make hybrid IT simple, to be the IT of industrial IoT powering the intelligent edge, and having the services expertise to make it happen.

In the end, it's all about the right mix of technology plus the right customer experience, driving the right outcomes.

Best Regards,

Antonio

Antonio Neri

Hewlett Packard Enterprise

Executive Vice President and General Manager, Enterprise Group

Introduction

Why we wrote this book

Workloads are key to every technology decision. Given today's unprecedented growth in IT consumption models to drive business success, more than ever before, the effective deployment of IT resources is critical. Making the right workload placement decisions is pivotal to delivering workloads-based business services in a way that is efficient, flexible, and scalable. To boil it down to a single question, every business operates a series of workloads and wants to understand, what are the ideal platforms to run those workloads on? We wrote this book to help CTOs, data center administrators, systems architects, IT professionals and other technical decision-makers discover some truly practical answers.

How this book is organized

This book covers a sampling of the most common workloads found in companies looking to evolve their data center and IT strategies. It seeks to present the detailed technical information that IT professionals require to identify the right workload platforms to meet their technical needs and achieve their business goals. Every workload has unique characteristics that require specific design features in order to run in an optimized manner. Workloads vary widely, such as scale-out versus scale-up, and have different characteristics, so understanding the fundamental aspects of the workload and crafting a solution that optimizes the operation of the workload is the key to high performance.

These workloads represent a variety of solutions designed by the authors of their respective chapters. The following is a list of the key topics with a short description of each chapter:

Anatomy of a Workload. Chapter 1 is written by a workload expert who provides a broad-based introduction to workloads from an industry perspective including categories of workloads, the data explosion, workload framework, and other high-level related topics.

Scale-Out. Chapters 2-4 cover this topic from three different angles: High Performance Computing (HPC) for actuarial modeling; HPC for scientific research; and scale out over many servers to improve the IT infrastructure.

Object storage and large file storage. Chapter 5 looks at using object storage software to archive petabytes of data across a distributed system, while Chapter 6 focuses on large files stored on disk using custom backup software. Each of these chapters include two case studies and their design outcomes.

Desktop Solutions: Chapters 7-10 cover four different workloads: Hosted Desktop Infrastructure (HDI) in which HPE Moonshot cartridges are devoted to a desktop; Virtual Desktop Infrastructure

(VDI); engineering Virtual Desktop Infrastructure (eVDI); and one chapter that focuses on Citrix. Many firms that had been testing various desktop solutions are now deploying such environments on a large scale and these four chapters provide a good sampling of solutions.

Electronic Medical Records (EMR): Chapter 11 shows an EMR workload solution using HPE ConvergedSystem 700 which is key to the strategy in many healthcare environments.

Hyper Converged System on Virtual Machines. Chapter 12 shows the design and implementation of an HPE HC 250 system to support a variety of virtualized applications.

Scale-up: Chapter 13 covers a SAP HANA solution on an HPE Superdome X; and Chapter 14 covers Microsoft SQL Server 2016 on Superdome X. Scale-up applications require features such as large memory capability one of the many benefits of Superdome X covered in these chapters.

Backup and Recovery: Chapters 15-17 cover three different backup and recovery-related solutions: Chapter 14 shows a way to deploy backup and recovery of virtual environments with Veeam and the requisite backup hardware; Chapter 15 covers the implementation of a Hyper Converged system with HPE StoreVirtual; while Chapter 16 details a Disaster Recovery (DR) solution that focuses on the storage aspects of DR.

Hadoop: Chapter 18 covers an asymmetric approach whereby compute and storage can be scaled independently of one another.

The last three chapters of the book focus on the workloads of the future. They do not cover a specific workload example but instead are high-level topics that highlight strategic choices for optimizing IT operations, including insights into future technologies that will apply to many workloads.

Hybrid IT: Chapter 19 looks at the background and strategic choices for identifying the “Right Mix” of applications to run in your private cloud and public cloud when creating and migrating to a Hybrid IT environment.

Internet of Things (IoT): Chapter 20 introduces IoT at a high-level to give a window into the business opportunities in this fast-growing technical area.

Composable Infrastructure: Chapter 21 covers advanced technologies such as HPE Synergy that is available today. Synergy provides a fluid pool of IT resources suited for many workloads and should be evaluated for any environment.

What you will learn

As you can see from the list of chapters, there are a wide variety of workloads covered as well as an overview of related topics such as IoT and Hybrid IT. This book will help define requirements and make the right choice of tools and technologies. Since each chapter focuses on a particular workload, it is designed for you to reference your specific areas of interest.

The layout of the chapters is similar in that most cover key aspects of the workload including a solution overview, the components used in the solution, the key considerations used to craft the solution,

and so on. Most of the chapters include a Bill of Materials (BOM) so that you can see the physical components included in the design. The BOMs are indicative of the type of workload being featured. A detailed customer needs assessment with a professional IT consultant is recommended to evaluate the exact BOM required in a specific instance.

This is a “blue collar” book in that that chapters don’t provide a lot of background on the workload topic but instead delve into the solution quickly. The BOM and rack diagrams give a clear vision of the solution implemented and some of the nuances related to the design. In some chapters, for instance, a wiring diagram is covered because it relates to the specific manner in which connections are employed in order to achieve high availability.

Some of the solutions were crafted over months and many iterations of the solution may have been needed to arrive at finished design. In many cases, this was a long and arduous process, that has been distilled in this book down to some essential ingredients. In the end, however, a solid design was produced and that is what you are seeing summarized in each chapter. We hope that this information will help you select the best platforms and leap the hurdles of designing and implementing powerful and effective workload solutions.

About the Authors

This book consists of submissions from many authors who crafted the solution covered in their respective chapter. I worked closely with the authors to produce a book that comes as close as possible to reading as if one author produced the entire book.

Primary Author

Marty Poniatowski is a Senior Director at Hewlett Packard Enterprise in the Business Development, Enablement, Solutions, and Technology (BEST) organization. He manages the presales technical experts in the U.S. East for many disciplines that represent the portfolio of HPE services and products. His focus is on the strategic initiatives of hybrid IT, intelligent edge, data-driven results, and many related technologies.

During his career Marty has authored 18 books on IT topics. Marty holds graduate degrees in Information Technology (New York University) and Management, and an undergraduate degree in Electrical Engineering.

Guest Authors

Mike Brito is a Cloud and Software-Defined Data Center (SDDC) technical sales team leader at HPE. Mike focuses on leadership, team building, creating business value through technical solutions, helping customers discover new opportunities, improving efficiency, and customer engagement. He is a graduate of Roger Williams University and Columbia University.

Steve Haldeman is currently Director of the HPE Americas Strategic Solutions Architecture team providing expert, technical presales solution services across key industry workload and solution areas including the Internet-of-Things (IoT). As a 35 year veteran with HPE, Steve had held many senior technical contributor roles including those in software development, project management, as Solution Architect for Unix systems, advanced systems consolidation, internet-based commerce and high-availability/mission critical system design and implementation. In addition, he has led teams of Solution Architect teams in many of these areas. Steve is an active HPE representative in the area of IoT to industry groups and with customers

John Tsang leads the Americas workload business strategy and go to market to unify a common cross functional approach to workloads & solutions. Leveraging the collective resources of the presales organization and worldwide and regional business units, John has been executing new PAN HPE sales plays and initiatives around the priority workloads.

Most recently John has taken on expanded leadership role to drive the IoT business for North America sales. His responsibilities include leading the IoT go-to-market activities which includes sales, marketing and operations readiness to build the IoT business. John has held numerous leadership roles in end user and channel sales, marketing and business development and joined HPE through the 3Com acquisition.

John graduated with a Bachelor of Business Administration from Loyola University of Chicago and continues to expand his education in coaching and leadership.

Chapter Authors

The following is a list of authors in alphabetical order:

Gary Allard is a Senior Solution Architect at HPE. Gary specializes in scale-out solutions designed for state government, local government, higher education and healthcare institutions.

Jeff Dabal is a senior Pre-Sales Channel Architect covering the enterprise solutions and services portfolio at HPE. Jeff works closely with HPE partners to provide enablement and related activities.

Michael Fahey is a Solution Architect at HPE focused on commercial and enterprise accounts. He specializes in storage technology and solutions.

Vikram Fernandes is a Senior Solution Architect at HPE focused on strategic customer accounts specializing in application and performance characterization.

Kevin Gildea is a Solution Architect at HPE focused on the Service Provider accounts. He specializes in distributed systems and scale-out design. Kevin contributed to multiple chapters in this book.

Rob Goldstein is an Enterprise Architect at HPE, where he supports some of HPE's largest customers by helping them solve challenging business problems through technology.

Anton Hagens is an Enterprise Architect at HPE focused on Commercial and Strategic Accounts with a focus on high-performance computing that meet the clients business outcomes.

Beth O'Malley is a Solution Architect at HPE. Her focus is Channel partner technical enablement that represents the portfolio of HPE services and products focused related to storage.

Zvadia Hibshoosh is a Solution Architect at HPE focused primarily on Industry Standard Servers (ISS) solutions in the Financial Services Industry (FSI).

George Riemer is an Enterprise Architect at HPE. George represents the full portfolio of products and services working directly with HPE enterprise customers.

Taras Rudko is a Presales Solution Architect in the Northeast for Hewlett Packard Enterprise specializing in storage workloads in virtualized environments.

Mary Alice Sallah is an HPE Solution Architect, working primarily with enterprise accounts in the Northeast to design workload-optimized compute infrastructure. Mary Alice submitted two chapters in this book.

Chuck Strobel is Master Solution Architect for HPE's Mission Critical Servers. Chuck is focused on HPE's scale-up Intel-based Servers, SAP HANA Converged Systems, and Integrity HP-UX Server solutions.

Brenda Suárez-Marill is a Solution Architect at HPE. She currently supports some of HPE's largest customers by helping them solve their most challenging technical problems to satisfy business requirements.

Tom Wilson is a Senior Infrastructure Architect for HPE's Enterprise Solutions and Architecture team. Tom focuses on Client Virtualization solutions including Virtual Desktop Infrastructure (VDI) and Server Based Computing (SBC).

Eric Wise is a HPE Master ASE focused on HPE's scale-up Intel-based Servers, SAP HANA Converged Systems, and Integrity HP-UX server solutions.

Benjamin Wold is a Storage Solution Architect at HPE. He specializes on storage workloads in virtualized environments.

Joseph Yanushpolsky is an Enterprise Solutions Architect with HPE Northeast team specializing in design, sizing, and tuning of complex infrastructure and application environments. Joseph contributed portions of two chapters in this book.

1 Anatomy of Workloads

INTRODUCTION

Establishing and running a top performing business requires not only a reliable but also a dynamic information and operations technology infrastructure to handle the demands of all the functional groups within an organization. The focus and commitment from the C-suite within a business to drive differentiated customer outcomes while doing so with higher operational efficiencies puts considerable demands on the information technology organization. With what we are now calling the Idea Economy, there has never been a time in our history where you can easily turn an idea into a business model and do it faster than ever before. This in turn can also bring more challenges to businesses where traditional competitors may not be the major threat anymore, but rather new upstart organizations that are created overnight and cause tremendous disruption.

This era of digital everything is transforming how information technology is consumed with an even a bigger focus around driving business outcomes. In order to achieve these outcomes, information technology solutions must be mindful of the key transformation areas of cloud computing with hybrid infrastructure, big data, mobility including workplace productivity, and security that have provided a perfect storm for this wholesale digital revolution. From cloud and/or multi-cloud computing to even bigger data with the Internet of Things (IoT), more modern approaches to handle and make meaning of this vast amount of structured and unstructured data is key. With the digital collaboration in the mobile workplace and the expansion of what we call the Intelligent Edge, having the levels of security to protect all the applications being accessed from the edge to data center requires intelligence and dynamic allocation of resources across the infrastructure layers.

At the same time, it is not enough to have a purpose-built appliance anymore with server or storage solutions handling one or two workloads. Composable and converged hardware platforms are now required to handle a diverse set of workloads and also have enough intelligence to recognize which workloads are prioritized over others to maximize efficiencies and productivity. Software-defined architectures enable the handling of the vast amount of workloads that have a tremendous influence on the creation of hybrid IT environments that we are adopting today.

Depending on whether the audience is an engineer, IT professional, or business leader, the terminology and taxonomy of what a workload is can vary. The basic definition of workload in computing terms is the amount of processing that the computer has been given to complete in a certain time and space. The workload consists of some level of application programming running on the computer and usually some number of users connected to and interacting with the computer's applications. According to TechTarget in 2006, a defined workload can be specified as a benchmark when evaluating a computer system in terms of performance (how easily the computer handles the workload), which in turn

is generally divided into response time (the time between a user request and a response to the request from the system) and throughput (how much work is accomplished over a period of time).

The evolution of workloads

As computing continues to evolve from centralized mainframe technologies in the 1970s and 1980s to distributed computing with client server technologies in the 1990s; to the 2000s' expansion of the Internet and with mobility currently driving cloud computing in the 2010s, how we move, access, and process workloads become front-and-center. As the next decade approaches, the explosion of IoT with increasing sensor technologies being embedded in everyday objects, the speed of how we handle the massive amount of data generated through IoT to make real-time actionable decisions through various applications and platforms securely will bring together the lines of information technology and operation technology like never before.

At its core, IoT is just another differentiated workload and, from a user perspective, it is all about gaining access immediately to information and applications from any device and being able to do something with it anywhere and at any time. This progression is depicted in Figure 1-1.

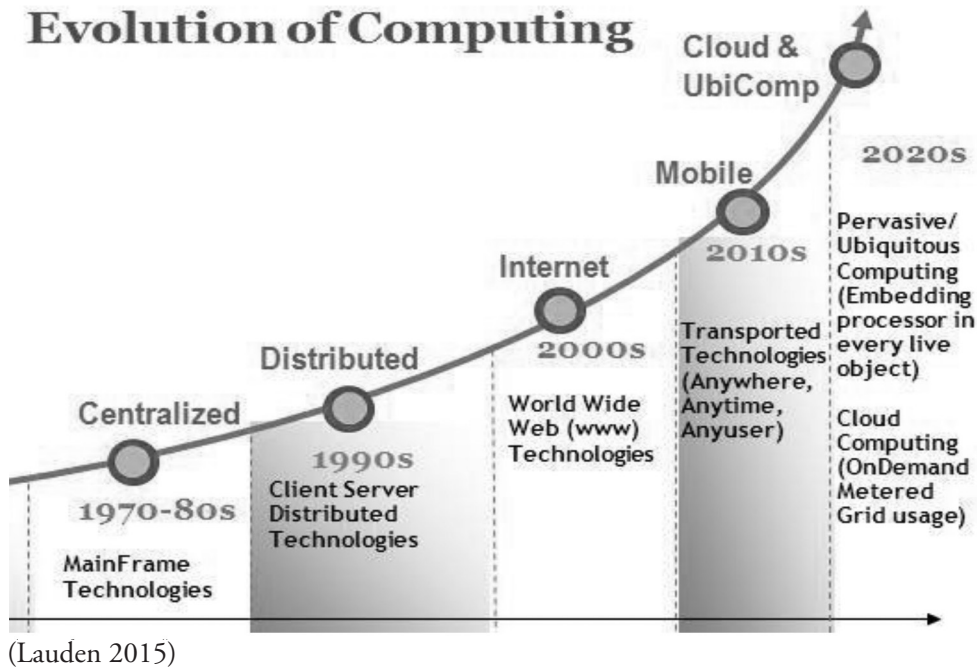


Figure 1-1 Evolution of computing

There is continued proliferation of mobile devices and growing use of new mobile collaboration tools. This is taking place in industries such as retail, financial services, manufacturing, and will continue to transition certain workloads and applications from on-premises to off-premises computing. The need to address workloads from a hybrid cloud and infrastructure perspective will continue to grow. Enabling on-premises private and hybrid clouds for workloads will be essential as customers prefer to have a hybrid environment to provide more flexibility and security for users. To better understand the shift of workloads to cloud environments, Figure 1-2 defines the various compute and multi-cloud models.

| Compute Environment | Definition |
|--|--|
| On-premises dedicated | The server infrastructure exists on-premises, is dedicated to and optimized for a single application, may or may not be virtualized, but is not cloud-based. |
| On-premises private cloud | The cloud infrastructure exists on-premises and is provisioned for exclusive use by a single company or institution. It is owned, managed, and operated by the company or institution. |
| Off-premises hosted or managed private cloud | The cloud infrastructure exists off-premises and is provisioned for exclusive use by a single company or institution. It may be owned, managed, and operated by the company, a third party, or some combination of them. |
| Public cloud | The cloud infrastructure is provisioned for use by multiple companies or institutions. The cloud infrastructure is owned, managed and operated by the cloud provider at its site. |
| Hybrid cloud | The cloud infrastructure is a composition of both private and public cloud infrastructures that remain unique entities, but are bound together by technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). |

(Source: IDC, 2016)

Figure 1-2 Cloud models

The challenge for many organizations will be to determine which workloads should run in what environment. Even more importantly determining the workload portability and suitability with the infrastructure options that can enable the various multi-cloud options and different consumption models. Recent research conducted by the TBR group in Figure 1-3 illustrates that more customers tend to perceive on-premises private cloud as more effective to scalability and having more ability to secure and integrate.