

NIST Special Publication 1500-6

**DRAFT NIST Big Data Interoperability
Framework:
Volume 6, Reference Architecture**

NIST Big Data Public Working Group
Reference Architecture Subgroup

Draft Version 1
April 6, 2015

<http://dx.doi.org/10.6028/NIST.SP.1500-6>

NIST
National Institute of
Standards and Technology
U.S. Department of Commerce

NIST Special Publication 1500-6
Information Technology Laboratory

**DRAFT NIST Big Data Interoperability
Framework:
Volume 6, Reference Architecture
Draft Version 1**

NIST Big Data Public Working Group (NBD-PWG)
Reference Architecture Subgroup
National Institute of Standards and Technology
Gaithersburg, MD 20899

April 2015



U. S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Dr. Willie E. May, Under Secretary of Commerce for Standards and Technology and Director

National Institute of Standards and Technology Special Publication 1500-6
60 pages (April 6, 2015)

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by Federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, Federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. All NIST Information Technology Laboratory publications, other than the ones noted above, are available at <http://www.nist.gov/publication-portal.cfm>.

Public comment period: April 6, 2015 through May 21, 2015

Comments on this publication may be submitted to Wo Chang

National Institute of Standards and Technology
Attn: Wo Chang, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8900) Gaithersburg, MD 20899-8930
Email: SP1500comments@nist.gov

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at NIST promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in Federal information systems. This document reports on ITL's research, guidance, and outreach efforts in Information Technology and its collaborative activities with industry, government, and academic organizations.

Abstract

Big Data is a term used to describe the large amount of data in the networked, digitized, sensor-laden, information-driven world. While opportunities exist with Big Data, the data can overwhelm traditional technical approaches and the growth of data is outpacing scientific and technological advances in data analytics. To advance progress in Big Data, the NIST Big Data Public Working Group (NBD-PWG) is working to develop consensus on important, fundamental concepts related to Big Data. The results are reported in the *NIST Big Data Interoperability Framework* series of volumes. This volume, Volume 6, summarizes the work performed by the NBD-PWG to characterize Big Data from an architecture perspective, presents the NIST Big Data Reference Architecture (NBDRA) conceptual model, and discusses the components and fabrics of the NBDRA.

Keywords

Big Data, reference architecture, System Orchestrator, Data Provider, Application Provider, Framework Provider, Data Consumer, Security and Privacy Fabric, Management Fabric, use cases, Big Data characteristics

Acknowledgements

This document reflects the contributions and discussions by the membership of the NBD-PWG, co-chaired by Wo Chang of the NIST ITL, Robert Marcus of ET-Strategies, and Chaitanya Baru, University of California San Diego Supercomputer Center.

The document contains input from members of the NBD-PWG: Reference Architecture Subgroup, led by Orit Levin (Microsoft), Don Krapohl (Augmented Intelligence), and James Ketner (AT&T); Technology Roadmap Subgroup, led by Carl Buffington (Vistronix), David Boyd (InCadence Strategic Solutions), and Dan McClary (Oracle); Definitions and Taxonomies Subgroup, led by Nancy Grady (SAIC), Natasha Balac (SDSC), and Eugene Luster (R2AD); Use Cases and Requirements Subgroup, led by Geoffrey Fox (University of Indiana) and Tsegereda Beyene(Cisco); Security and Privacy Subgroup, led by Arnab Roy (Fujitsu) and Akhil Manchanda (GE).

NIST SP1500-6, Version 1 has been collaboratively authored by the NBD-PWG. As of the date of publication, there are over six hundred NBD-PWG participants from industry, academia, and government. Federal agency participants include the National Archives and Records Administration (NARA), National Aeronautics and Space Administration (NASA), National Science Foundation (NSF), and the U.S. Departments of Agriculture, Commerce, Defense, Energy, Health and Human Services, Homeland Security, Transportation, Treasury, and Veterans Affairs.

NIST acknowledges the specific contributions to this volume by the following NBD-PWG members:

Chaitan Baru <i>University of California, San Diego, Supercomputer Center</i>	Pavithra Kenjige <i>PK Technologies</i>	Felix Njeh <i>U.S. Department of the Army</i>
Janis Beach <i>Information Management Services, Inc.</i>	James Kobielus <i>IBM</i>	Gururaj Pandurangi <i>Aryan Consulting Corp.</i>
David Boyd <i>InCadence Strategic Solutions</i>	Donald Krapohl <i>Augmented Intelligence</i>	Linda Pelekoudas <i>Strategy and Design Solutions</i>
Scott Brim <i>Internet2</i>	Orit Levin <i>Microsoft</i>	Dave Raddatz <i>Silicon Graphics International Corp.</i>
Gregg Brown <i>Microsoft</i>	Eugene Luster <i>DISA/R2AD</i>	John Rogers <i>HP</i>
Carl Buffington <i>Vistronix</i>	Serge Manning <i>Huawei USA</i>	Arnab Roy <i>Fujitsu</i>
Yuri Demchenko <i>University of Amsterdam</i>	Robert Marcus <i>ET-Strategies</i>	Michael Seablom <i>NASA</i>
Jill Gemmill <i>Clemson University</i>	Gary Mazzaferro <i>AlloyCloud, Inc.</i>	Rupinder Singh <i>McAfee, Inc.</i>
Nancy Grady <i>SAIC</i>	Shawn Miller <i>U.S. Department of Veterans Affairs</i>	Anil Srivastava <i>Open Health Systems Laboratory</i>
Ronald Hale <i>ISACA</i>	Sanjay Mishra <i>Verizon</i>	Glenn Wasson <i>SAIC</i>
Keith Hare <i>JCC Consulting, Inc.</i>	Vivek Navale <i>NARA</i>	Timothy Zimmerlin <i>Automation Technologies Inc.</i>
Richard Jones <i>The Joseki Group LLC</i>	Quy Nguyen <i>NARA</i>	Alicia Zuniga-Alvarado <i>Consultant</i>

The editors for this document were Orit Levin, David Boyd, and Wo Chang.

Notice to Readers

NIST is seeking feedback on the proposed working draft of the *NIST Big Data Interoperability Framework: Volume 6, Reference Architecture*. Once public comments are received, compiled, and addressed by the NBD-PWG, and reviewed and approved by NIST internal editorial board, Version 1 of this volume will be published as final. Three versions are planned for this volume, with versions 2 and 3 building on the first. Further explanation of the three planned versions and the information contained therein is included in Section 1.5 of this document.

Please be as specific as possible in any comments or edits to the text. Specific edits include, but are not limited to, changes in the current text, additional text further explaining a topic or explaining a new topic, additional references, or comments about the text, topics, or document organization. These specific edits can be recorded using one of the two following methods.

1. **TRACK CHANGES**: make edits to and comments on the text directly into this Word document using track changes
2. **COMMENT TEMPLATE**: capture specific edits using the Comment Template (http://bigdatawg.nist.gov/_uploadfiles/SP1500-1-to-7_comment_template.docx), which includes space for Section number, page number, comment, and text edits

Submit the edited file from either method 1 or 2 to SP1500comments@nist.gov with the volume number in the subject line (e.g., Edits for Volume 6.)

Please contact Wo Chang (wchang@nist.gov) with any questions about the feedback submission process.

Big Data professionals continue to be welcome to join the NBD-PWG to help craft the work contained in the volumes of the NIST Big Data Interoperability Framework. Additional information about the NBD-PWG can be found at <http://bigdatawg.nist.gov>.

Table of Contents

EXECUTIVE SUMMARY	VII
1 INTRODUCTION	1
1.1 BACKGROUND	1
1.2 SCOPE AND OBJECTIVES OF THE REFERENCE ARCHITECTURES SUBGROUP	2
1.3 REPORT PRODUCTION	3
1.4 REPORT STRUCTURE	3
1.5 FUTURE WORK ON THIS VOLUME	4
2 HIGH LEVEL REFERENCE ARCHITECTURE REQUIREMENTS.....	5
2.1 USE CASES AND REQUIREMENTS	5
2.2 REFERENCE ARCHITECTURE SURVEY	7
2.3 TAXONOMY	7
3 NBDRA CONCEPTUAL MODEL.....	10
4 FUNCTIONAL COMPONENTS OF THE NBDRA	12
4.1 SYSTEM ORCHESTRATOR	12
4.2 DATA PROVIDER	12
4.3 BIG DATA APPLICATION PROVIDER	14
4.3.1 <i>Collection</i>	15
4.3.2 <i>Preparation</i>	15
4.3.3 <i>Analytics</i>	15
4.3.4 <i>Visualization</i>	16
4.3.5 <i>Access</i>	16
4.4 BIG DATA FRAMEWORK PROVIDER	16
4.4.1 <i>Infrastructure Frameworks</i>	17
4.4.2 <i>Data Platform Frameworks</i>	19
4.4.3 <i>Processing Frameworks</i>	23
4.4.4 <i>Messaging/Communications Frameworks</i>	28
4.4.5 <i>Resource Management Framework</i>	28
4.5 DATA CONSUMER	29
5 MANAGEMENT FABRIC OF THE NBDRA	31
5.1 SYSTEM MANAGEMENT	31
5.2 BIG DATA LIFECYCLE MANAGEMENT.....	32
6 SECURITY AND PRIVACY FABRIC OF THE NBDRA	34
7 CONCLUSION	35
APPENDIX A: DEPLOYMENT CONSIDERATIONS	A-1
APPENDIX B: TERMS AND DEFINITIONS	B-1
APPENDIX C: EXAMPLES OF BIG DATA ORGANIZATION APPROACHES	C-1
APPENDIX D: ACRONYMS.....	D-1
APPENDIX E: RESOURCES AND REFERENCES	E-1

Figures

FIGURE 1: NBDRA TAXONOMY	8
FIGURE 2: NIST BIG DATA REFERENCE ARCHITECTURE (NBDRA).....	10
FIGURE 3: DATA ORGANIZATION APPROACHES.....	20
FIGURE 4: DATA STORAGE TECHNOLOGIES	23
FIGURE 5: INFORMATION FLOW	24
FIGURE A-1: BIG DATA FRAMEWORK DEPLOYMENT OPTIONS.....	A-1
FIGURE B-1: DIFFERENCES BETWEEN ROW ORIENTED AND COLUMN ORIENTED STORES	B-3
FIGURE B-2: COLUMN FAMILY SEGMENTATION OF THE COLUMNAR STORES MODEL	B-3
FIGURE B-3: OBJECT NODES AND RELATIONSHIPS OF GRAPH DATABASES.....	B-6

Tables

TABLE 1: MAPPING USE CASE CHARACTERIZATION CATEGORIES TO REFERENCE ARCHITECTURE COMPONENTS AND FABRICS	5
TABLE 2: 13 DWARFS—ALGORITHMS FOR SIMULATION IN THE PHYSICAL SCIENCES.....	25

DRAFT

1 Executive Summary

2 The NIST Big Data Public Working group (NBD-PWG) Reference Architecture Subgroup prepared this
3 *NIST Big Data Interoperability Framework: Volume 6, Reference Architecture* to provide a vendor-
4 neutral, technology- and infrastructure-agnostic conceptual model and examine related issues. The
5 conceptual model, referred to as the NIST Big Data Reference Architecture (NBDRA), was crafted by
6 examining publicly available Big Data architectures representing various approaches and products. Inputs
7 from the other NBD-PWG subgroups were also incorporated into the creation of the NBDRA. It is
8 applicable to a variety of business environments, including tightly-integrated enterprise systems, as well
9 as loosely-coupled vertical industries that rely on cooperation among independent stakeholders. The
10 NBDRA captures the two known Big Data economic value chains: information, where value is created by
11 data collection, integration, analysis, and applying the results to data-driven services; and, the information
12 technology (IT), where value is created by providing networking, infrastructure, platforms, and tools in
13 support of vertical, data-based applications.

14 The *NIST Big Data Interoperability Framework* consists of seven volumes, each of which addresses a
15 specific key topic, resulting from the work of the NBD-PWG. The seven volumes are as follows:

- 16 • Volume 1, Definitions
- 17 • Volume 2, Taxonomies
- 18 • Volume 3, Use Cases and General Requirements
- 19 • Volume 4, Security and Privacy
- 20 • Volume 5, Architectures White Paper Survey
- 21 • Volume 6, Reference Architecture
- 22 • Volume 7, Standards Roadmap

23 The *NIST Big Data Interoperability Framework* will be released in three versions, which correspond to
24 the three stages of the NBD-PWG work. The three stages aim to achieve the following with respect to the
25 NBDRA:

- 26 Stage 1: Identify the high-level Big Data reference architecture key components, which are
27 technology, infrastructure, and vendor agnostic
- 28 Stage 2: Define general interfaces between the NBDRA components
- 29 Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces

30 Potential areas of future work for the Subgroup during stage 2 are highlighted in Section 1.5 of this
31 volume. The current effort documented in this volume reflects concepts developed within the rapidly
32 evolving field of Big Data.

33

1 INTRODUCTION

1.1 BACKGROUND

There is broad agreement among commercial, academic, and government leaders about the remarkable potential of Big Data to spark innovation, fuel commerce, and drive progress. Big Data is the common term used to describe the deluge of data in today's networked, digitized, sensor-laden, and information-driven world. The availability of vast data resources carries the potential to answer questions previously out of reach, including the following:

- How can a potential pandemic reliably be detected early enough to intervene?
- Can new materials with advanced properties be predicted before these materials have ever been synthesized?
- How can the current advantage of the attacker over the defender in guarding against cyber-security threats be reversed?

There is also broad agreement on the ability of Big Data to overwhelm traditional approaches. The growth rates for data volumes, speeds, and complexity are outpacing scientific and technological advances in data analytics, management, transport, and data user spheres.

Despite widespread agreement on the inherent opportunities and current limitations of Big Data, a lack of consensus on some important, fundamental questions continues to confuse potential users and stymie progress. These questions include the following:

- What attributes define Big Data solutions?
- How is Big Data different from traditional data environments and related applications?
- What are the essential characteristics of Big Data environments?
- How do these environments integrate with currently deployed architectures?
- What are the central scientific, technological, and standardization challenges that need to be addressed to accelerate the deployment of robust Big Data solutions?

Within this context, on March 29, 2012, the White House announced the Big Data Research and Development Initiative.¹ The initiative's goals include helping to accelerate the pace of discovery in science and engineering, strengthening national security, and transforming teaching and learning by improving the ability to extract knowledge and insights from large and complex collections of digital data.

Six federal departments and their agencies announced more than \$200 million in commitments spread across more than 80 projects, which aim to significantly improve the tools and techniques needed to access, organize, and draw conclusions from huge volumes of digital data. The initiative also challenged industry, research universities, and nonprofits to join with the federal government to make the most of the opportunities created by Big Data.

Motivated by the White House initiative and public suggestions, the National Institute of Standards and Technology (NIST) has accepted the challenge to stimulate collaboration among industry professionals to further the secure and effective adoption of Big Data. As one result of NIST's Cloud and Big Data Forum held on January 15–17, 2013, there was strong encouragement for NIST to create a public working group for the development of a Big Data Interoperability Framework. Forum participants noted that this roadmap should define and prioritize Big Data requirements, including interoperability, portability, reusability, extensibility, data usage, analytics, and technology infrastructure. In doing so, the roadmap would accelerate the adoption of the most secure and effective Big Data techniques and technology.

76 On June 19, 2013, the NIST Big Data Public Working Group (NBD-PWG) was launched with extensive
 77 participation by industry, academia, and government from across the nation. The scope of the NBD-PWG
 78 involves forming a community of interests from all sectors—including industry, academia, and
 79 government—with the goal of developing consensus on definitions, taxonomies, secure reference
 80 architectures, security and privacy, and—from these—a standards roadmap. Such a consensus would
 81 create a vendor-neutral, technology- and infrastructure-independent framework that would enable Big
 82 Data stakeholders to identify and use the best analytics tools for their processing and visualization
 83 requirements on the most suitable computing platform and cluster, while also allowing value-added from
 84 Big Data service providers.

85 The *NIST Big Data Interoperability Framework* consists of seven volumes, each of which addresses a
 86 specific key topic, resulting from the work of the NBD-PWG. The seven volumes are as follows:

- 87 • Volume 1, Definitions
- 88 • Volume 2, Taxonomies
- 89 • Volume 3, Use Cases and General Requirements
- 90 • Volume 4, Security and Privacy
- 91 • Volume 5, Architectures White Paper Survey
- 92 • Volume 6, Reference Architecture
- 93 • Volume 7, Standards Roadmap

94 The *NIST Big Data Interoperability Framework* will be released in three versions, which correspond to
 95 the three stages of the NBD-PWG work. The three stages aim to achieve the following with respect to the
 96 NBDRA:

- 97 Stage 1: Identify the high-level Big Data reference architecture key components, which are
 98 technology, infrastructure, and vendor agnostic
- 99 Stage 2: Define general interfaces between the NBDRA components
- 100 Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces

101 Potential areas of future work for the Subgroup during stage 2 are highlighted in Section 1.5 of this
 102 volume. The current effort documented in this volume reflects concepts developed within the rapidly
 103 evolving field of Big Data.

104 **1.2 SCOPE AND OBJECTIVES OF THE REFERENCE ARCHITECTURES SUBGROUP**

105 Reference architectures provide “an authoritative source of information about a specific subject area that
 106 guides and constrains the instantiations of multiple architectures and solutions.”² Reference architectures
 107 generally serve as a foundation for solution architectures and may also be used for comparison and
 108 alignment of instantiations of architectures and solutions.

109 The goal of the NBD-PWG Reference Architecture Subgroup is to develop an open reference architecture
 110 for Big Data that achieves the following objectives:

- 111 • Provides a common language for the various stakeholders
- 112 • Encourages adherence to common standards, specifications, and patterns
- 113 • Provides consistent methods for implementation of technology to solve similar problem sets
- 114 • Illustrates and improves understanding of the various Big Data components, processes, and
 115 systems, in the context of a vendor- and technology- agnostic Big Data conceptual model
- 116 • Provides a technical reference for U.S. government departments, agencies, and other consumers
 117 to understand, discuss, categorize, and compare Big Data solutions
- 118 • Facilitates analysis of candidate standards for interoperability, portability, reusability, and
 119 extensibility

120 The NIST Big Data Reference Architecture (NBDRA) is a high-level conceptual model crafted to serve
 121 as a tool to facilitate open discussion of the requirements, design structures, and operations inherent in
 122 Big Data. The NBDRA is intended to facilitate the understanding of the operational intricacies in Big
 123 Data. It does not represent the system architecture of a specific Big Data system, but rather is a tool for
 124 describing, discussing, and developing system-specific architectures using a common framework of
 125 reference. The model is not tied to any specific vendor products, services, or reference implementation,
 126 nor does it define prescriptive solutions that inhibit innovation.

127 The NBDRA does not address the following:

- 128 • Detailed specifications for any organization’s operational systems
- 129 • Detailed specifications of information exchanges or services
- 130 • Recommendations or standards for integration of infrastructure products

131 **1.3 REPORT PRODUCTION**

132 A wide spectrum of Big Data architectures have been explored and developed as part of various industry,
 133 academic, and government initiatives. The development of the NBDRA and material contained in this
 134 volume involved the following steps:

- 135 1. Announce that the NBD-PWG Reference Architecture Subgroup is open to the public to attract
 136 and solicit a wide array of subject matter experts and stakeholders in government, industry, and
 137 academia
- 138 2. Gather publicly-available Big Data architectures and materials representing various stakeholders,
 139 different data types, and diverse use cases ^a
- 140 3. Examine and analyze the Big Data material to better understand existing concepts, usage, goals,
 141 objectives, characteristics, and key elements of Big Data, and then document the findings using
 142 NIST’s Big Data taxonomies model (presented in *NIST Big Data Interoperability Framework:
 143 Volume 2, Taxonomies*)
- 144 4. Develop a technology-independent, open reference architecture based on the analysis of Big Data
 145 material and inputs received from other NBD-PWG subgroups

146 **1.4 REPORT STRUCTURE**

147 The organization of this document roughly corresponds to the process used by the NBD-PWG to develop
 148 the NBDRA. Following the introductory material presented in Section 1, the remainder of this document
 149 is organized as follows:

- 150 • Section 2 contains high-level, system requirements in support of Big Data relevant to the design
 151 of the NBDRA and discusses the development of these requirements
- 152 • Section 3 presents the generic, technology-independent NBDRA conceptual model
- 153 • Section 4 discusses the five main functional components of the NBDRA
- 154 • Section 5 describes the system and lifecycle management considerations related to the NBDRA
 155 management fabric
- 156 • Section 6 briefly introduces security and privacy topics related to the security and privacy fabric
 157 of the NBDRA
- 158 • Appendix A summarizes deployment considerations
- 159 • Appendix B lists the terms and definitions in this document
- 160 • Appendix C provides examples of Big Data logical data architecture options

^a Many of the architecture use cases were originally collected by the NBD-PWG Use Case and Requirements Subgroup and can be accessed at <http://bigdatawg.nist.gov/usecases.php>.

- 161 • Appendix D defines the acronyms used in this document
162 • Appendix E lists several general resources that provide additional information on topics covered
163 in this document and specific references in this document

164 **1.5 FUTURE WORK ON THIS VOLUME**

165 The *NIST Big Data Interoperability Framework* will be released in three versions, which correspond to
166 the three stages of the NBD-PWG work. The three stages aim to achieve the following with respect to the
167 NBDRA:

- 168 Stage 1: Identify the common reference architecture components of Big Data implementations and
169 formulate the technology-independent NBDRA
170 Stage 2: Define general interfaces between the NBDRA components
171 Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces

172 This document (Version 1) presents the overall NBDRA components and fabrics with high-level
173 description and functionalities.

174 Version 2 activities will focus on the definition of general interfaces between the NBDRA components by
175 performing the following:

- 176 • Select use cases from the 62 (51 general and 11 security and privacy) submitted use cases or
177 other, to be identified, meaningful use cases
- 178 • Work with domain experts to identify workflow and interactions among the NBDRA
179 components and fabrics
- 180 • Explore and model these interactions within a small scale, manageable, and well-defined
181 confined environment
- 182 • Aggregate the common data workflow and interactions between NBDRA components and
183 fabrics and package them into general interfaces

184 Version 3 activities will focus on validation of the NBDRA through the use of the defined NBDRA
185 general interfaces to build general Big Data applications. The validation strategy will include the
186 following:

- 187 • Implement the same set of use cases used in Version 2 by using the defined general interfaces
- 188 • Identify and implement a few new use cases outside the Version 2 scenarios
- 189 • Enhance general NBDRA interfaces through lessons learned from the implementations in
190 Version 3 activities

191 The general interfaces developed during Version 2 activities will offer a starting point for further
192 refinement by any interested parties and is not intended to be a definitive solution to address all
193 implementation needs.

194

195

196 2 HIGH LEVEL REFERENCE ARCHITECTURE REQUIREMENTS

197 The development of a Big Data reference architecture requires a thorough understanding of current
 198 techniques, issues, and concerns. To this end, the NBD-PWG collected use cases to gain an understanding
 199 of current applications of Big Data, conducted a survey of reference architectures to understand
 200 commonalities within Big Data architectures in use, developed a taxonomy to understand and organize
 201 the information collected, and reviewed existing technologies and trends relevant to Big Data. The results
 202 of these NBD-PWG activities were used in the development of the NBDRA and are briefly described in
 203 this section.

204 2.1 USE CASES AND REQUIREMENTS

205 To develop the use cases, publically available information was collected for various Big Data
 206 architectures in nine broad areas, or application domains. Participants in the NBD-PWG Use Case and
 207 Requirements Subgroup and other interested parties provided the use case details via a template, which
 208 helped standardize the responses and facilitate subsequent analysis and comparison of the use cases.
 209 However, submissions still varied in levels of detail, quantitative data, or qualitative information. The
 210 *NIST Big Data Interoperability Framework: Volume 3, Use Cases and General Requirements* document
 211 presents the original use cases, an analysis of the compiled information, and the requirements extracted
 212 from the use cases.

213 The extracted requirements represent challenges faced in seven characterization categories (Table 1)
 214 developed by the Subgroup. Requirements specific to the use cases were aggregated into high level,
 215 generalized requirements, which are vendor and technology neutral.

216 The use case characterization categories were used as input in the development of the NBDRA and map
 217 directly to NBDRA components and fabrics as shown in Table 1.

218 *Table 1: Mapping Use Case Characterization Categories to Reference Architecture Components and*
 219 *Fabrics*

USE CASE CHARACTERIZATION CATEGORIES		REFERENCE ARCHITECTURE COMPONENTS AND FABRICS
Data sources	→	Data Provider
Data transformation	→	Big Data Application Provider
Capabilities	→	Big Data Framework Provider
Data consumer	→	Data Consumer
Security and privacy	→	Security and Privacy Fabric
Lifecycle management	→	System Orchestrator; Management Fabric
Other requirements	→	To all components and fabrics

220
 221 The high-level, generalized requirements are presented below. The development of these generalized
 222 requirements is presented in the *NIST Big Data Interoperability Framework: Volume 3, Use Cases and*
 223 *Requirements* document.

224 **DATA PROVIDER REQUIREMENTS**

- 225 • DSR-1: Reliable, real-time, asynchronous, streaming, and batch processing to collect data from
- 226 centralized, distributed, and cloud data sources, sensors, or instruments
- 227 • DSR-2: Slow, bursty, and high throughput data transmission between data sources and
- 228 computing clusters
- 229 • DSR-3: Diversified data content ranging from structured and unstructured text, documents,
- 230 graphs, web sites, geospatial, compressed, timed, spatial, multimedia, simulation, and
- 231 instrumental (i.e., system managements and monitoring) data

232 **BIG DATA APPLICATION PROVIDER REQUIREMENTS**

- 233 • TPR-1: Diversified, compute-intensive, analytic processing and machine learning techniques
- 234 • TPR-2: Batch and real time analytic processing
- 235 • TPR-3: Processing large diversified data content and modeling
- 236 • TPR-4: Processing data in motion (e.g., streaming, fetching new content, data tracking,
- 237 traceability, data change management, and data boundaries)

238 **BIG DATA FRAMEWORK PROVIDER REQUIREMENTS**

- 239 • CPR-1: Legacy software and advanced software packages
- 240 • CPR-2: Legacy and advanced computing platforms
- 241 • CPR-3: Legacy and advanced distributed computing clusters, co-processors, input/output (I/O)
- 242 processing
- 243 • CPR-4: Advanced networks (e.g., Software Defined Networks [SDN]) and elastic data
- 244 transmission, including fiber, cable, and wireless networks (e.g., local area network [LAN], wide
- 245 area network [WAN], metropolitan area network [MAN], Wi-Fi)
- 246 • CPR-5: Legacy, large, virtual, and advanced distributed data storage
- 247 • CPR-6: Legacy and advanced programming executables, applications, tools, utilities, and
- 248 libraries

249 **DATA CONSUMER REQUIREMENTS**

- 250 • DCR-1: Fast searches (~0.1 seconds) from processed data with high relevancy, accuracy, and
- 251 recall
- 252 • DCR-2: Diversified output file formats for visualization, rendering, and reporting
- 253 • DCR-3: Visual layout for results presentation
- 254 • DCR-4: Rich user interface for access using browser, visualization tools
- 255 • DCR-5: High resolution, multi-dimension layer of data visualization
- 256 • DCR-6: Streaming results to clients

257 **SECURITY AND PRIVACY REQUIREMENTS**

- 258 • SPR-1: Protect and preserve security and privacy of sensitive data
- 259 • SPR-2: Support sandbox, access control, and multi-tenant, multi-level, policy-driven
- 260 authentication on protected data and ensure these are in line with accepted governance, risk, and
- 261 compliance (GRC) and confidentiality, integrity, and availability (CIA) best practices

262 **MANAGEMENT REQUIREMENTS**

- 263 • LMR-1: Data quality curation, including pre-processing, data clustering, classification,
- 264 reduction, and format transformation
- 265 • LMR-2: Dynamic updates on data, user profiles, and links
- 266 • LMR-3: Data lifecycle and long-term preservation policy, including data provenance
- 267 • LMR-4: Data validation

- 268 • LMR-5: Human annotation for data validation
- 269 • LMR-6: Prevention of data loss or corruption
- 270 • LMR-7: Multi-site (including cross-border, geographically dispersed) archives
- 271 • LMR-8: Persistent identifier and data traceability
- 272 • LMR-9: Standardization, aggregation, and normalization of data from disparate sources

273 **OTHER REQUIREMENTS**

- 274 • OR-1: Rich user interface from mobile platforms to access processed results
- 275 • OR-2: Performance monitoring on analytic processing from mobile platforms
- 276 • OR-3: Rich visual content search and rendering from mobile platforms
- 277 • OR-4: Mobile device data acquisition and management
- 278 • OR-5: Security across mobile devices and other smart devices such as sensors

279 **2.2 REFERENCE ARCHITECTURE SURVEY**

280 The NBD-PWG Reference Architecture Subgroup conducted a survey of current reference architectures
 281 to advance the understanding of the operational intricacies in Big Data and to serve as a tool for
 282 developing system-specific architectures using a common reference framework. The Subgroup surveyed
 283 currently published Big Data platforms by leading companies or individuals supporting the Big Data
 284 framework and analyzed the collected material. This effort revealed a consistency between Big Data
 285 architectures that served in the development of the NBDRA. Survey details, methodology, and
 286 conclusions are reported in *NIST Big Data Interoperability Framework: Volume 5, Architectures White
 287 Paper Survey*.

288 **2.3 TAXONOMY**

289 The NBD-PWG Definitions and Taxonomy Subgroup focused on identifying Big Data concepts, defining
 290 terms needed to describe the new Big Data paradigm, and defining reference architecture terms. The
 291 reference architecture taxonomy presented below provides a hierarchy of the components of the reference
 292 architecture. Additional taxonomy details are presented in the *NIST Big Data Interoperability
 293 Framework: Volume 2, Taxonomy* document.

294 Figure 1 outlines potential actors for the seven roles developed by the NBD-PWG Definition and
 295 Taxonomy Subgroup. The blue boxes contain the name of the role at the top with potential actors listed
 296 directly below.

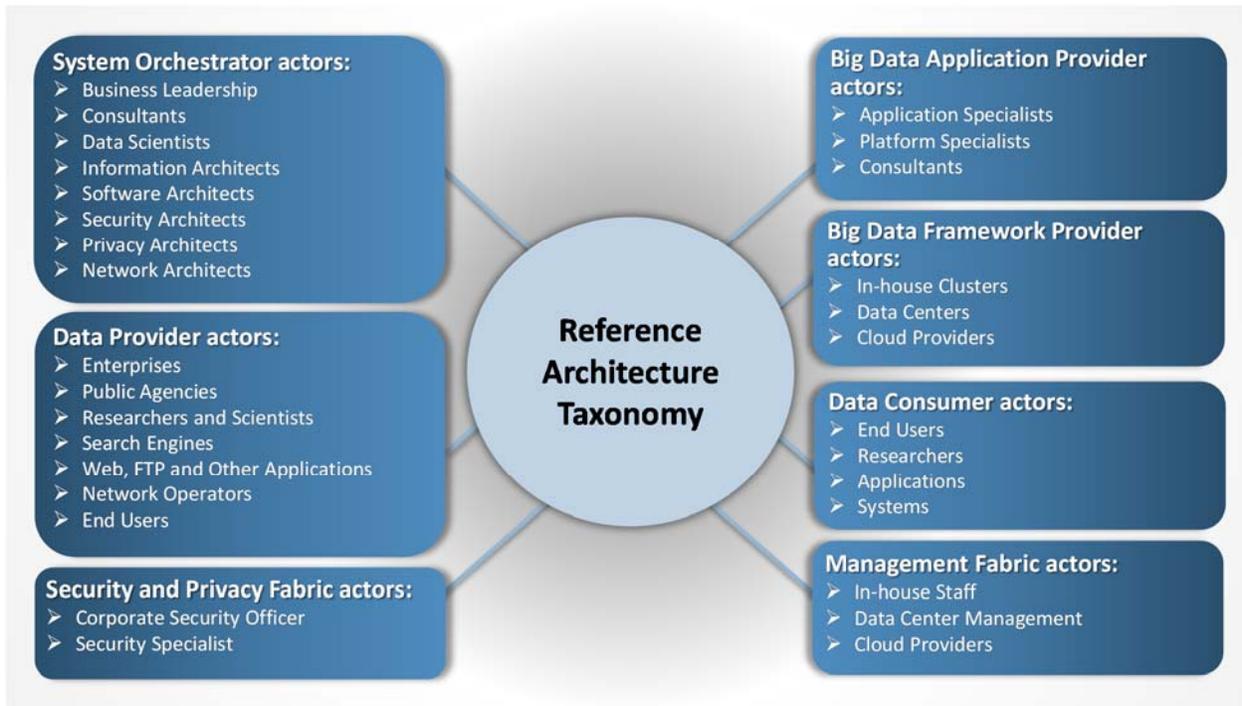


Figure 1: NBDRA Taxonomy

297

298 **SYSTEM ORCHESTRATOR**

299 The System Orchestrator provides the overarching requirements that the system must fulfill, including
 300 policy, governance, architecture, resources, and business requirements, as well as monitoring or auditing
 301 activities to ensure the system complies with those requirements. The System Orchestrator role provides
 302 system requirements, high-level design, and monitoring for the data system. While the role pre-dates Big
 303 Data systems, some related design activities have changed within the Big Data paradigm.

304 **DATA PROVIDER**

305 A Data Provider makes data available to itself or to others. In fulfilling its role, the Data Provider creates
 306 an abstraction of various types of data sources (such as raw data or data previously transformed by
 307 another system) and makes them available through different functional interfaces. The actor fulfilling this
 308 role can be part of the Big Data system, internal to the organization in another system, or external to the
 309 organization orchestrating the system. While the concept of a Data Provider is not new, the greater data
 310 collection and analytics capabilities have opened up new possibilities for providing valuable data.

311 **BIG DATA APPLICATION PROVIDER**

312 The Big Data Application Provider executes the manipulations of the data lifecycle to meet requirements
 313 established by the System Orchestrator. This is where the general capabilities within the Big Data
 314 framework are combined to produce the specific data system. While the activities of an application
 315 provider are the same whether the solution being built concerns Big Data or not, the methods and
 316 techniques have changed because the data and data processing is parallelized across resources.

317 **BIG DATA FRAMEWORK PROVIDER**

318 The Big Data Framework Provider has general resources or services to be used by the Big Data
 319 Application Provider in the creation of the specific application. There are many new components from
 320 which the Big Data Application Provider can choose in using these resources and the network to build the
 321 specific system. This is the role that has seen the most significant changes because of Big Data. The Big
 322 Data Framework Provider consists of one or more instances of the three subcomponents: infrastructure

323 frameworks, data platforms, and processing frameworks. There is no requirement that all instances at a
324 given level in the hierarchy be of the same technology and, in fact, most Big Data implementations are
325 hybrids combining multiple technology approaches. These provide flexibility and can meet the complete
326 range of requirements that are driven from the Big Data Application Provider. Due to the rapid emergence
327 of new techniques, this is an area that will continue to need discussion.

328 **DATA CONSUMER**

329 The Data Consumer receives the value output of the Big Data system. In many respects, it is the recipient
330 of the same type of functional interfaces that the Data Provider exposes to the Big Data Application
331 Provider. After the system adds value to the original data sources, the Big Data Application Provider then
332 exposes that same type of functional interfaces to the Data Consumer.

333 **SECURITY AND PRIVACY FABRIC**

334 Security and privacy issues affect all other components of the NBDRA. The Security and Privacy Fabric
335 interacts with the System Orchestrator for policy, requirements, and auditing and also with both the Big
336 Data Application Provider and the Big Data Framework Provider for development, deployment, and
337 operation. The *NIST Big Data Interoperability Framework: Volume 4, Security and Privacy* document
338 discusses security and privacy topics.

339 **MANAGEMENT FABRIC**

340 The Big Data characteristics of volume, velocity, variety, and variability demand a versatile system and
341 software management platform for provisioning, software and package configuration and management,
342 along with resource and performance monitoring and management. Big Data management involves
343 system, data, security, and privacy considerations at scale, while maintaining a high level of data quality
344 and secure accessibility.

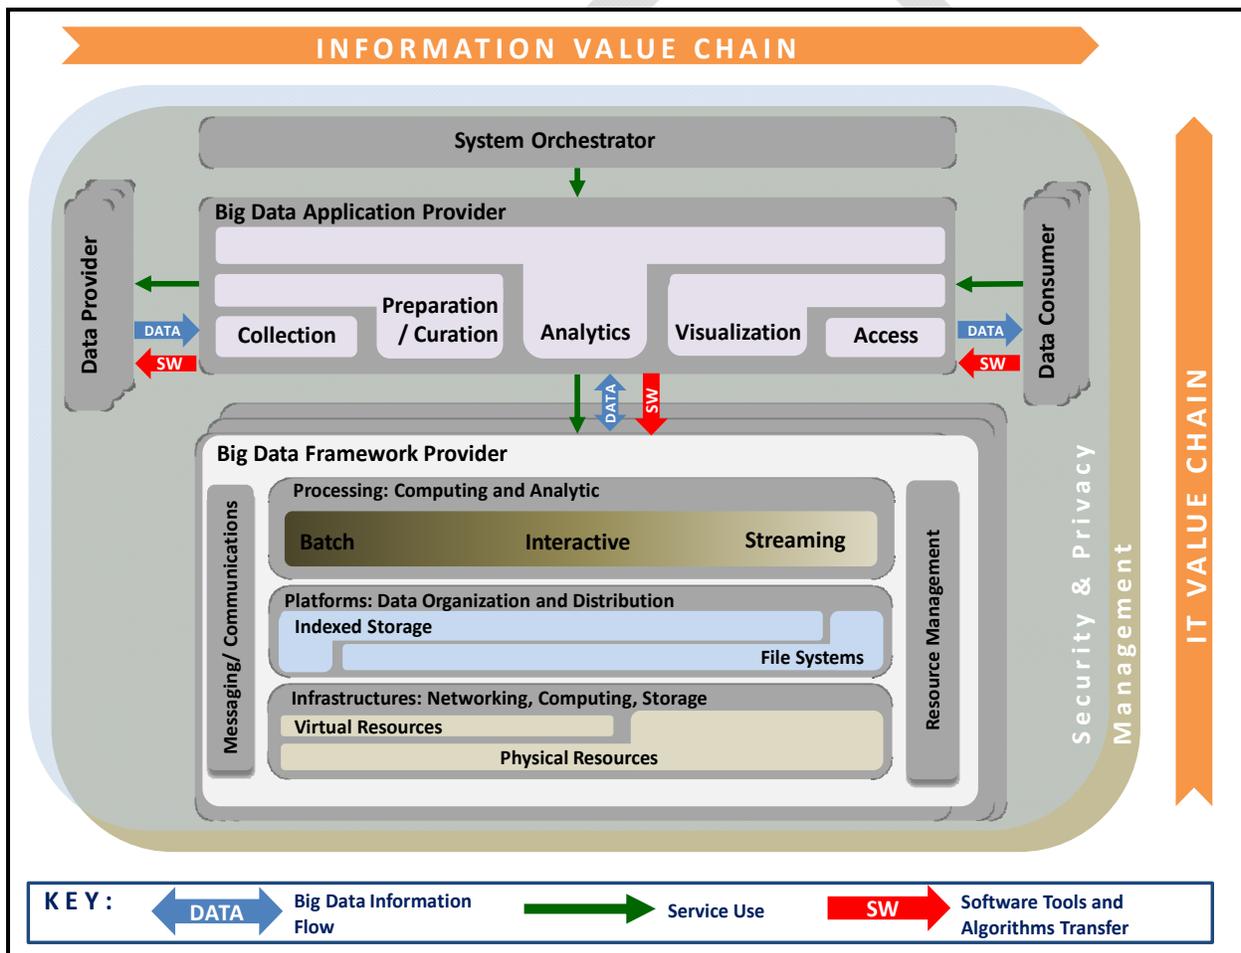
345

346

347 3 NBDRA CONCEPTUAL MODEL

348 As discussed in Section 2, the NBD-PWG Reference Architecture Subgroup used a variety of inputs from
 349 other NBD-PWG subgroups in developing a vendor-neutral, technology- and infrastructure-agnostic,
 350 conceptual model of Big Data architecture. This conceptual model, the NBDRA, is shown in Figure 2 and
 351 represents a Big Data system comprised of five logical functional components connected by
 352 interoperability interfaces (i.e., services). Two fabrics envelop the components, representing the
 353 interwoven nature of management and security and privacy with all five of the components.

354 The NBDRA is intended to enable system engineers, data scientists, software developers, data architects,
 355 and senior decision makers to develop solutions to issues that require diverse approaches due to
 356 convergence of Big Data characteristics within an interoperable Big Data ecosystem. It provides a
 357 framework to support a variety of business environments, including tightly-integrated enterprise systems
 358 and loosely-coupled vertical industries, by enhancing understanding of how Big Data complements and
 359 differs from existing analytics, business intelligence, databases, and systems.



360

361 *Figure 2: NIST Big Data Reference Architecture (NBDRA)*

362 The NBDRA is organized around two axes representing the two Big Data value chains: the information
 363 (horizontal axis) and the IT (vertical axis). Along the information axis, the value is created by data

364 collection, integration, analysis, and applying the results following the value chain. Along the IT axis, the
 365 value is created by providing networking, infrastructure, platforms, application tools, and other IT
 366 services for hosting of and operating the Big Data in support of required data applications. At the
 367 intersection of both axes is the Big Data Application Provider component, indicating that data analytics
 368 and its implementation provide the value to Big Data stakeholders in both value chains. The names of the
 369 Big Data Application Provider and Big Data Framework Provider components contain “providers” to
 370 indicate that these components provide or implement a specific technical function within the system.

371 The five main NBDRA components, shown in Figure 2 and discussed in detail in Section 4, represent
 372 different technical roles that exist in every Big Data system. These functional components are as follows:

- 373 • System Orchestrator
- 374 • Data Provider
- 375 • Big Data Application Provider
- 376 • Big Data Framework Provider
- 377 • Data Consumer

378 The two fabrics shown in Figure 2 encompassing the five functional components are the following:

- 379 • Management
- 380 • Security and Privacy

381 These two fabrics provide services and functionality to the five functional components in the areas
 382 specific to Big Data and are crucial to any Big Data solution.

383 The “DATA” arrows in Figure 2 show the flow of data between the system’s main components. Data
 384 flows between the components either physically (i.e., by value) or by providing its location and the means
 385 to access it (i.e., by reference). The “SW” arrows show transfer of software tools for processing of Big
 386 Data *in situ*. The “Service Use” arrows represent software programmable interfaces. While the main focus
 387 of the NBDRA is to represent the run-time environment, all three types of communications or transactions
 388 can happen in the configuration phase as well. Manual agreements (e.g., service-level agreements
 389 [SLAs]) and human interactions that may exist throughout the system are not shown in the NBDRA.

390 The components represent functional roles in the Big Data ecosystem. In system development, actors and
 391 roles have the same relationship as in the movies, but system development actors can represent
 392 individuals, organizations, software, or hardware. According to the Big Data taxonomy, a single actor can
 393 play multiple roles, and multiple actors can play the same role. The NBDRA does not specify the business
 394 boundaries between the participating actors or stakeholders, so the roles can either reside within the same
 395 business entity or can be implemented by different business entities. Therefore, the NBDRA is applicable
 396 to a variety of business environments, from tightly-integrated enterprise systems to loosely-coupled
 397 vertical industries that rely on the cooperation of independent stakeholders. As a result, the notion of
 398 internal versus external functional components or roles does not apply to the NBDRA. However, for a
 399 specific use case, once the roles are associated with specific business stakeholders, the functional
 400 components would be considered as internal or external—subject to the use case’s point of view.

401 The NBDRA does support the representation of stacking or chaining of Big Data systems. For example, a
 402 Data Consumer of one system could serve as a Data Provider to the next system down the stack or chain.

403

404 4 FUNCTIONAL COMPONENTS OF THE NBDRA

405 As outlined in Section 3, the five main functional components of the NBDRA represent the different
 406 technical roles within a Big Data system. The functional components are listed below and discussed in
 407 subsequent subsections.

- 408 • **System Orchestrator:** Defines and integrates the required data application activities into an
 409 operational vertical system
- 410 • **Data Provider:** Introduces new data or information feeds into the Big Data system
- 411 • **Big Data Application Provider:** Executes a data lifecycle to meet security and privacy
 412 requirements as well as System Orchestrator-defined requirements
- 413 • **Big Data Framework Provider:** Establishes a computing framework in which to execute
 414 certain transformation applications while protecting the privacy and integrity of data
- 415 • **Data Consumer:** Includes end users or other systems that use the results of the Big Data
 416 Application Provider

417 4.1 SYSTEM ORCHESTRATOR

418 The System Orchestrator role includes defining and integrating the required data application activities
 419 into an operational vertical system. Typically, the System Orchestrator involves a collection of more
 420 specific roles, performed by one or more actors, which manage and orchestrate the operation of the Big
 421 Data system. These actors may be human components, software components, or some combination of the
 422 two. The function of the System Orchestrator is to configure and manage the other components of the Big
 423 Data architecture to implement one or more workloads that the architecture is designed to execute. The
 424 workloads managed by the System Orchestrator may be assigning/provisioning framework components to
 425 individual physical or virtual nodes at the lower level, or providing a graphical user interface that supports
 426 the specification of workflows linking together multiple applications and components at the higher level.
 427 The System Orchestrator may also, through the Management Fabric, monitor the workloads and system to
 428 confirm that specific quality of service requirements are met for each workload, and may actually
 429 elastically assign and provision additional physical or virtual resources to meet workload requirements
 430 resulting from changes/surges in the data or number of users/transactions.

431 The NBDRA represents a broad range of Big Data systems, from tightly coupled enterprise solutions
 432 (integrated by standard or proprietary interfaces) to loosely coupled vertical systems maintained by a
 433 variety of stakeholders bounded by agreements and standard or standard-de-facto interfaces.

434 In an enterprise environment, the System Orchestrator role is typically centralized and can be mapped to
 435 the traditional role of system governor that provides the overarching requirements and constraints, which
 436 the system must fulfill, including policy, architecture, resources, or business requirements. A system
 437 governor works with a collection of other roles (e.g., data manager, data security, and system manager) to
 438 implement the requirements and the system's functionality.

439 In a loosely coupled vertical system, the System Orchestrator role is typically decentralized. Each
 440 independent stakeholder is responsible for its own system management, security, and integration, as well
 441 as integration within the Big Data distributed system using the interfaces provided by other stakeholders.

442 4.2 DATA PROVIDER

443 The Data Provider role introduces new data or information feeds into the Big Data system for discovery,
 444 access, and transformation by the Big Data system. New data feeds are distinct from the data already in
 445 use by the system and residing in the various system repositories. Similar technologies can be used to

446 access both new data feeds and existing data. The Data Provider actors can be anything from a sensor, to
447 a human inputting data manually, to another Big Data system.

448 One of the important characteristics of a Big Data system is the ability to import and use data from a
449 variety of data sources. Data sources can be internal or public records, tapes, images, audio, videos,
450 sensor data, web logs, system and audit logs, HTTP cookies, and other sources. Humans, machines,
451 sensors, online and offline applications, Internet technologies, and other actors can also produce data
452 sources. The roles of Data Provider and Big Data Application Provider often belong to different
453 organizations, unless the organization implementing the Big Data Application Provider owns the data
454 sources. Consequently, data from different sources may have different security and privacy
455 considerations. In fulfilling its role, the Data Provider creates an abstraction of the data sources. In the
456 case of raw data sources, the Data Provider can potentially cleanse, correct, and store the data in an
457 internal format that is accessible to the Big Data system that will ingest it.

458 The Data Provider can also provide an abstraction of data previously transformed by another system (i.e.,
459 legacy system, another Big Data system.) In this case, the Data Provider would represent a Data
460 Consumer of the other system. For example, Data Provider 1 could generate a streaming data source from
461 the operations performed by Data Provider 2 on a dataset at rest.

462 Data Provider activities include the following, which are common to most systems that handle data:

- 463 • Collecting the data
- 464 • Persisting the data
- 465 • Providing transformation functions for data scrubbing of sensitive information such as
466 Personally Identifiable Information (PII)
- 467 • Creating the metadata describing the data source(s), usage policies/access rights, and other
468 relevant attributes
- 469 • Enforcing access rights on data access
- 470 • Establishing formal or informal contracts for data access authorizations
- 471 • Making the data accessible through suitable programmable push or pull interfaces
- 472 • Providing push or pull access mechanisms
- 473 • Publishing the availability of the information and the means to access it

474 The Data Provider exposes a collection of interfaces (or services) for discovering and accessing the data.
475 These interfaces would typically include a registry so that applications can locate a Data Provider,
476 identify the data of interest it contains, understand the types of access allowed, understand the types of
477 analysis supported, locate the data source, determine data access methods, identify the data security
478 requirements, identify the data privacy requirements, and other pertinent information. Therefore, the
479 interface would provide the means to register the data source, query the registry, and identify a standard
480 set of data contained by the registry.

481 Subject to Big Data characteristics (i.e., volume, variety, velocity, and variability) and system design
482 considerations, interfaces for exposing and accessing data would vary in their complexity and can include
483 both push and pull software mechanisms. These mechanisms can include subscription to events, listening
484 to data feeds, querying for specific data properties or content, and the ability to submit a code for
485 execution to process the data *in situ*. Because the data can be too large to economically move across the
486 network, the interface could also allow the submission of analysis requests (e.g., software code
487 implementing a certain algorithm for execution), with the results returned to the requestor. Data access
488 may not always be automated, but might involve a human role logging into the system and providing
489 directions where new data should be transferred (e.g., establishing a subscription to an email based data
490 feed).

491 The interface between the Data Provider and Big Data Application Provider typically will go through
492 three phases: initiation, data transfer, and termination. The initiation phase is started by either party and

493 often includes some level of authentication/authorization. The phase may also include queries for
 494 metadata about the source or consumer, such as the list of available topics in a publish/subscribe
 495 (pub/sub) model and the transfer of any parameters (e.g., object count/size limits or target storage
 496 locations). Alternatively, the phase may be as simple as one side opening a socket connection to a known
 497 port on the other side.

498 The data transfer phase may be a push from the Data Provider or a pull by the Big Data Application
 499 Provider. It may also be a singular transfer or involve multiple repeating transfers. In a repeating transfer
 500 situation, the data may be a continuous stream of transactions/records/bytes. In a push scenario, the Big
 501 Data Application Provider must be prepared to accept the data asynchronously but may also be required
 502 to acknowledge (or negatively acknowledge) the receipt of each unit of data. In a pull scenario, the Big
 503 Data Application Provider would specifically generate a request that defines through parameters of the
 504 data to be returned. The returned data could itself be a stream or multiple records/units of data and the
 505 data transfer phase may consist of multiple request/send transactions.

506 The termination phase could be as simple as one side simply dropping the connection or could include
 507 checksums, counts, hashes, or other information about the completed transfer.

508 **4.3 BIG DATA APPLICATION PROVIDER**

509 The Big Data Application Provider role executes a specific set of operations along the data lifecycle to
 510 meet the requirements established by the System Orchestrator, as well as meeting security and privacy
 511 requirements. The Big Data Application Provider is the architecture component that encapsulates the
 512 business logic and functionality to be executed by the architecture. The Big Data Application Provider
 513 activities include the following

- 514 • Collection
- 515 • Preparation
- 516 • Analytics
- 517 • Visualization
- 518 • Access

519 These activities are represented by the subcomponents of the Big Data Application Provider as shown in
 520 Figure 2. The execution of these activities would typically be specific to the application and, therefore,
 521 are not candidates for standardization. However, the metadata and the policies defined and exchanged
 522 between the application's subcomponents could be standardized when the application is specific to a
 523 vertical industry.

524 While many of these activities exist in traditional data processing systems, the data volume, velocity,
 525 variety, and variability present in Big Data systems radically change their implementation. Traditional
 526 algorithms and mechanisms of traditional data processing implementations need to be adjusted and
 527 optimized to create applications that are responsive and can grow to handle ever growing data collections.

528 As data propagates through the ecosystem, it is being processed and transformed in different ways in
 529 order to extract the value from the information. Each activity of the Big Data Application Provider can be
 530 implemented by independent stakeholders and deployed as stand-alone services.

531 The Big Data Application Provider can be a single instance or a collection of more granular Big Data
 532 Application Providers, each implementing different steps in the data lifecycle. Each of the activities of the
 533 Big Data Application Provider may be a general service invoked by the System Orchestrator, Data
 534 Provider, or Data Consumer, such as a Web server, a file server, a collection of one or more application
 535 programs, or a combination. There may be multiple and differing instances of each activity or a single
 536 program may perform multiple activities. Each of the activities is able to interact with the underlying Big
 537 Data Framework Providers as well as with the Data Providers and Data Consumers. In addition, these

538 activities may execute in parallel or in any number of sequences and will frequently communicate with
 539 each other through the messaging/communications element of the Big Data Framework Provider. Also,
 540 the functions of the Big Data Application Provider, specifically the collection and access activities, will
 541 interact with the Security and Privacy Fabric to perform authentication/authorization and record/maintain
 542 data provenance.

543 Each of the functions can run on a separate Big Data Framework Provider or all can use a common Big
 544 Data Framework Provider. The considerations behind these different system approaches would depend on
 545 potentially different technological needs, business and/or deployment constraints (including privacy), and
 546 other policy considerations. The baseline NBDRA does not show the underlying technologies, business
 547 considerations, and topological constraints, thus making it applicable to any kind of system approach and
 548 deployment.

549 For example, the infrastructure of the Big Data Application Provider would be represented as one of the
 550 Big Data Framework Providers. If the Big Data Application Provider uses external/outourced
 551 infrastructures as well, it or they will be represented as another or multiple Big Data Framework
 552 Providers in the NBDRA. The multiple grey blocks behind the Big Data Framework Providers in Figure 2
 553 indicate that multiple Big Data Framework Providers can support a single Big Data Application Provider.

554 **4.3.1 COLLECTION**

555 In general, the collection activity of the Big Data Application Provider handles the interface with the Data
 556 Provider. This may be a general service, such as a file server or web server configured by the System
 557 Orchestrator to accept or perform specific collections of data, or it may be an application specific service
 558 designed to pull data or receive pushes of data from the Data Provider. Since this activity is receiving data
 559 at a minimum, it must store/buffer the received data until it is persisted through the Big Data Framework
 560 Provider. This persistence need not be to physical media but may simply be to an in-memory queue or
 561 other service provided by the processing frameworks of the Big Data Framework Provider. The collection
 562 activity is likely where the extraction portion of the Extract, Transform, Load/Extract, Load, Transform
 563 (ETL/ELT) cycle is performed. At the initial collection stage, sets of data (e.g., data records) of similar
 564 structure are collected (and combined), resulting in uniform security, policy, and other considerations.
 565 Initial metadata is created (e.g., subjects with keys are identified) to facilitate subsequent aggregation or
 566 lookup methods.

567 **4.3.2 PREPARATION**

568 The preparation activity is where the transformation portion of the ETL/ELT cycle is likely performed,
 569 although analytics activity will also likely perform advanced parts of the transformation. Tasks performed
 570 by this activity could include data validation (e.g., checksums/hashes, format checks), cleansing (e.g.,
 571 eliminating bad records/fields), outlier removal, standardization, reformatting, or encapsulating. This
 572 activity is also where source data will frequently be persisted to archive storage in the Big Data
 573 Framework Provider and provenance data will be verified or attached/associated. Verification or
 574 attachment may include optimization of data through manipulations (e.g., deduplication) and indexing to
 575 optimize the analytics process. This activity may also aggregate data from different Data Providers,
 576 leveraging metadata keys to create an expanded and enhanced data set.

577 **4.3.3 ANALYTICS**

578 The analytics activity of the Big Data Application Provider includes the encoding of the low-level
 579 business logic of the Big Data system (with higher level business process logic being encoded by the
 580 System Orchestrator). The activity implements the techniques to extract knowledge from the data based
 581 on the requirements of the vertical application. The requirements specify the data processing algorithms
 582 for processing the data to produce new insights that will address the technical goal. The analytics activity
 583 will leverage the processing frameworks to implement the associated logic. This typically involves the
 584 activity providing software that implements the analytic logic to the batch and/or streaming elements of

585 the processing framework for execution. The messaging/communication framework of the Big Data
586 Framework Provider may be used to pass data or control functions to the application logic running in the
587 processing frameworks. The analytic logic may be broken up into multiple modules to be executed by the
588 processing frameworks which communicate, through the messaging/communication framework, with
589 each other and other functions instantiated by the Big Data Application Provider.

590 **4.3.4 VISUALIZATION**

591 The visualization activity of the Big Data Application Provider prepares elements of the processed data
592 and the output of the analytic activity for presentation to the Data Consumer. The objective of this activity
593 is to format and present data in such a way as to optimally communicate meaning and knowledge. The
594 visualization preparation may involve producing a text-based report or rendering the analytic results as
595 some form of graphic. The resulting output may be a static visualization and may simply be stored
596 through the Big Data Framework Provider for later access. However, the visualization activity frequently
597 interacts with the access activity, the analytics activity, and the Big Data Framework Provider (processing
598 and platform) to provide interactive visualization of the data to the Data Consumer based on parameters
599 provided to the access activity by the Data Consumer. The visualization activity may be completely
600 application implemented, leverage one or more application libraries, or may use specialized visualization
601 processing frameworks within the Big Data Framework Provider.

602 **4.3.5 ACCESS**

603 The access activity within the Big Data Application Provider is focused on the communication/interaction
604 with the Data Consumer. Similar to the collection activity, the access activity may be a generic service
605 such as a web server or application server that is configured by the System Orchestrator to handle specific
606 requests from the Data Consumer. This activity would interface with the visualization and analytic
607 activities to respond to requests from the Data Consumer (who may be a person) and uses the processing
608 and platform frameworks to retrieve data to respond to Data Consumer requests. In addition, the access
609 activity confirms that descriptive and administrative metadata and metadata schemes are captured and
610 maintained for access by the Data Consumer and as data is transferred to the Data Consumer. The
611 interface with the Data Consumer may be synchronous or asynchronous in nature and may use a pull or
612 push paradigm for data transfer.

613 **4.4 BIG DATA FRAMEWORK PROVIDER**

614 The Big Data Framework Provider typically consists of one or more hierarchically organized instances of
615 the components in the NBDRA IT value chain (Figure 2). There is no requirement that all instances at a
616 given level in the hierarchy be of the same technology. In fact, most Big Data implementations are
617 hybrids that combine multiple technology approaches in order to provide flexibility or meet the complete
618 range of requirements, which are driven from the Big Data Application Provider.

619 Many of the recent advances related to Big Data have been in the area of frameworks designed to scale to
620 Big Data needs (e.g., addressing volume, variety, velocity, and variability) while maintaining linear or
621 near linear performance. These advances have generated much of the technology excitement in the Big
622 Data space. Accordingly, there is a great deal more information available in the frameworks area
623 compared to the other components and the additional detail provided for the Big Data Framework
624 Provider in this document reflects this imbalance.

625 The Big Data Framework Provider is comprised of the following three subcomponents (from the bottom
626 to the top):

- 627 • Infrastructure Frameworks
- 628 • Data Platform Frameworks
- 629 • Processing Frameworks

630 **4.4.1 INFRASTRUCTURE FRAMEWORKS**

631 This Big Data Framework Provider element provides all of the resources necessary to host/run the
 632 activities of the other components of the Big Data system. Typically, these resources consist of some
 633 combination of physical resources, which may host/support similar virtual resources. These resources are
 634 generally classified as follows:

- 635 • Networking: These are the resources that transfer data from one infrastructure framework
 636 component to another
- 637 • Computing: These are the physical processors and memory that execute and hold the software of
 638 the other Big Data system components
- 639 • Storage: These are resources which provide persistence of the data in a Big Data system
- 640 • Environmental: These are the physical plant resources (power, cooling) that must be accounted
 641 for when establishing an instance of a Big Data system

642 While the Big Data Framework Provider component may be deployed directly on physical resources or
 643 on virtual resources, at some level all resources have a physical representation. Physical resources are
 644 frequently used to deploy multiple components that will be duplicated across a large number of physical
 645 nodes to provide what is known as horizontal scalability. Virtualization is frequently used to achieve
 646 elasticity and flexibility in the allocation of physical resources and is often referred to as Infrastructure as
 647 a Service (IaaS) within the cloud computing community. Virtualization is typically found in one of three
 648 basic forms within a Big Data Architecture.

- 649 • Native: In this form, a hypervisor runs natively on the bare metal and manages multiple virtual
 650 machines consisting of Operating Systems (OS) and applications.
- 651 • Hosted: In this form, an OS runs natively on the bare metal and a hypervisor runs on top of that to
 652 host a client OS and applications. This model is not often seen in Big Data architectures due to
 653 the increased overhead of the extra OS layer.
- 654 • Containerized: In this form, hypervisor functions are embedded in the OS, which runs on bare
 655 metal. Applications are run inside containers, which control or limit access to the OS and physical
 656 machine resources. This approach has gained popularity for Big Data architectures because it
 657 further reduces overhead since most OS functions are a single shared resource. It may not be
 658 considered as secure or stable since in the event that the container controls/limits fail, one
 659 application may take down every application sharing those physical resources.

660 The following subsections describe the types of physical and virtual resources that comprise Big Data
 661 infrastructure.

662 **4.4.1.1 NETWORKING**

663 The connectivity of the architecture infrastructure should be addressed, as it affects the velocity
 664 characteristic of Big Data. While, some Big Data implementations may solely deal with data that is
 665 already resident in the data center and does not need to leave the confines of the local network, others
 666 may need to plan and account for the movement of Big Data either into or out of the data center. The
 667 location of Big Data systems with transfer requirements may depend on the availability of external
 668 network connectivity (i.e., bandwidth) and the limitations of Transmission Control Protocol (TCP) where
 669 there is low latency (as measured by packet Round Trip Time) with the primary senders or receivers of
 670 Big Data. To address the limitations of TCP, architects for Big Data systems may need to consider some
 671 of the advanced non-TCP based communications protocols available that are specifically designed to
 672 transfer large files such as video and imagery.

673 Overall availability of the external links is another infrastructure aspect relating to the velocity
 674 characteristic of Big Data that should be considered in architecting external connectivity. A given
 675 connectivity link may be able to easily handle the velocity of data while operating correctly. However,
 676 should the quality of service on the link degrade or the link fail completely, data may be lost or simply

677 back up to the point that it can never recover. Use cases exist where the contingency planning for network
678 outages involves transferring data to physical media and physically transporting it to the desired
679 destination. However, even this approach is limited by the time it may require to transfer the data to
680 external media for transport.

681 The volume and velocity characteristics of Big Data often are driving factors in the implementation of the
682 internal network infrastructure as well. For example, if the implementation requires frequent transfers of
683 large multi-gigabyte files between cluster nodes, then high speed and low latency links are required.
684 Depending on the availability requirements, redundant and fault tolerant links may be required. Other
685 aspects of the network infrastructure include name resolution (e.g., Domain Name Server [DNS]) and
686 encryption along with firewalls and other perimeter access control capabilities. Finally, the network
687 infrastructure may also include automated deployment, provisioning capabilities, or agents and
688 infrastructure wide monitoring agents that are leveraged by the management/communication elements to
689 implement a specific model.

690 Two concepts, Software Defined Networks (SDN) and Network Function Virtualization (NFV), have
691 recently been developed in support of scalable networks and scalable systems using them.

692 **4.4.1.1 Software Defined Networks**

693 Frequently ignored, but critical to the performance of distributed systems and frameworks, and especially
694 critical to Big Data implementations, is the efficient and effective management of networking resources.
695 Significant advances in network resource management have been realized through what is known as
696 SDN. Much like virtualization frameworks manage shared pools of CPU/memory/disk, SDNs (or virtual
697 networks) manage pools of physical network resources. In contrast to the traditional approaches of
698 dedicated physical network links for data, management, I/O, and control, SDNs contain multiple physical
699 resources (including links and actual switching fabric) that are pooled and allocated as required to specific
700 functions and sometimes to specific applications. This allocation can consist of raw bandwidth, quality of
701 service (QOS) priority, and even actual data routes.

702 **4.4.1.2 Network Function Virtualization**

703 With the advent of virtualization, virtual appliances can now reasonably support a large number of
704 network functions that were traditionally performed by dedicated devices. Network functions that can be
705 implemented in this manner include routing/routers, perimeter defense (e.g., firewalls), remote access
706 authorization, and network traffic/load monitoring. Some key advantages of NFV include elasticity, fault
707 tolerance, and resource management. For example, the ability to automatically deploy/provision
708 additional firewalls in response to a surge in user or data connections and then un-deploy them when the
709 surge is over can be critical in handling the volumes associated with Big Data.

710 **4.4.1.2 COMPUTING**

711 The logical distribution of cluster/computing infrastructure may vary from a dense grid of physical
712 commodity machines in a rack, to a set of virtual machines running on a cloud service provider, or to a
713 loosely coupled set of machines distributed around the globe providing access to un-used computing
714 resources. Computing infrastructure also frequently includes the underlying operating systems and
715 associated services used to interconnect the cluster resources via the networking elements.

716 **4.4.1.3 STORAGE**

717 The storage infrastructure may include any resource from isolated local disks to Storage Area Networks
718 (SANs) or network attached storage.

719 Two aspects of storage infrastructure technology that directly influence their suitability for Big Data
720 solutions are capacity and transfer bandwidth. Capacity refers to the ability to handle the data volume.
721 Local disks/file systems are specifically limited by the size of the available media. Hardware or software
722 (HW/SW) redundant array of independent disks (RAID) solutions—in this case local to a processing

723 node—help with scaling by allowing multiple pieces of media to be treated as a single device. However,
724 this approach is limited by the physical dimension of the media and the number of devices the node can
725 accept. SAN and network-attached storage (NAS) implementations—often known as shared disk
726 solutions—remove that limit by consolidating storage into a storage specific device. By consolidating
727 storage, the second aspect—transfer bandwidth—may become an issue. While both network and I/O
728 interfaces are getting faster and many implementations support multiple transfer channels, I/O bandwidth
729 can still be a limiting factor. In addition, despite the redundancies provided by RAID, hot spares, multiple
730 power supplies, and multiple controllers, these boxes can often become I/O bottlenecks or single points of
731 failure in an enterprise. Many Big Data implementations address these issues by using distributed file
732 systems within the platform framework.

733 **4.4.1.4 ENVIRONMENTAL RESOURCES**

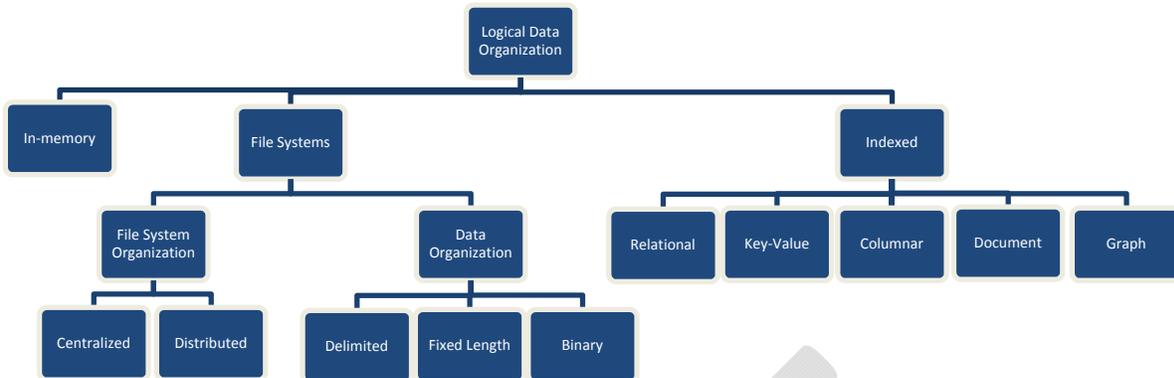
734 Environmental resources, such as power and heating, ventilation, and air conditioning (HVAC), are
735 critical to the Big Data Framework Provider. While environmental resources are critical to the operation
736 of the Big Data system, they are not within the technical boundaries and are, therefore, not depicted in
737 Figure 2, the NBDRA conceptual model.

738 Adequately sized infrastructure to support application requirements is critical to the success of Big Data
739 implementations. The infrastructure architecture operational requirements range from basic power and
740 cooling to external bandwidth connectivity (as discussed above). A key evolution that has been driven by
741 Big Data is the increase in server density (i.e., more CPU/memory/disk per rack unit). However, with this
742 increased density, infrastructure—specifically power and cooling—may not be distributed within the data
743 center to allow for sufficient power to each rack or adequate air flow to remove excess heat. In addition,
744 with the high cost of managing energy consumption within data centers, technologies have been
745 developed that actually power down or idle resources not in use to save energy or to reduce consumption
746 during peak periods.

747 **4.4.2 DATA PLATFORM FRAMEWORKS**

748 Data Platform Frameworks provide for the logical data organization and distribution combined with the
749 associated access application programming interfaces (APIs) or methods. The frameworks may also
750 include data registry and metadata services along with semantic data descriptions such as formal
751 ontologies or taxonomies. The logical data organization may range from simple delimited flat files to
752 fully distributed relational or columnar data stores. The storage mediums range from high latency robotic
753 tape drives, to spinning magnetic media, to flash/solid state disks, or to random access memory.
754 Accordingly, the access methods may range from file access APIs to query languages such as SQL.
755 Typical Big Data framework implementations would support either basic file system style storage or in-
756 memory storage and one or more indexed storage approaches. Based on the specific Big Data system
757 considerations, this logical organization may or may not be distributed across a cluster of computing
758 resources.

759 In most aspects, the logical data organization and distribution in Big Data storage frameworks mirrors the
760 common approach for most legacy systems. Figure 3 presents a brief overview of data organization
761 approaches for Big Data.



762
763 **Figure 3: Data Organization Approaches**

764 Many Big Data logical storage organizations leverage the common file system concept—where chunks of
765 data are organized into a hierarchical namespace of directories—as their base and then implement various
766 indexing methods within the individual files. This allows many of these approaches to be run both on
767 simple local storage file systems for testing purposes or on fully distributed file systems for scale.

768 **4.4.2.1 IN-MEMORY**

769 The infrastructure illustrated in the NBDRA (Figure 2) indicates that physical resources are required to
770 support analytics. However, such infrastructure will vary (i.e., will be optimized) for the Big Data
771 characteristics of the problem under study. Large, but static, historical datasets with no urgent analysis
772 time constraints would optimize the infrastructure for the volume characteristic of Big Data, while time-
773 critical analyses such as intrusion detection or social media trend analysis would optimize the
774 infrastructure for the velocity characteristic of Big Data. Velocity implies the necessity for extremely fast
775 analysis and the infrastructure to support it—namely, very low latency, in-memory analytics.

776 In-memory database technologies are increasingly used due to the significant reduction in memory prices
777 and the increased scalability of modern servers and operating systems. Yet, an in-memory element of a
778 velocity-oriented infrastructure will require more than simply massive random-access memory (RAM). It
779 will also require optimized data structures and memory access algorithms to fully exploit RAM
780 performance. Current in-memory database offerings are beginning to address this issue.

781 Traditional database management architectures are designed to use spinning disks as the primary storage
782 mechanism, with the main memory of the computing environment relegated to providing caching of data
783 and indexes. Many of these in-memory storage mechanisms have their roots in the massively parallel
784 processing and super computer environments popular in the scientific community.

785 These approaches should not be confused with solid state (e.g., flash) disks or tiered storage systems that
786 implement memory-based storage which simply replicate the disk style interfaces and data structures but
787 with faster storage medium. Actual in-memory storage systems typically eschew the overhead of file
788 system semantics and optimize the data storage structure to minimize memory footprint and maximize the
789 data access rates. These in-memory systems may implement general purpose relational and other NoSQL
790 style organization and interfaces or be completely optimized to a specific problem and data structure.

791 Like traditional disk-based systems for Big Data, these implementations frequently support horizontal
792 distribution of data and processing across multiple independent nodes—although shared memory
793 technologies are still prevalent in specialized implementations. Unlike traditional disk-based approaches,
794 in-memory solutions and the supported applications must account for the lack of persistence of the data
795 across system failures. Some implementations leverage a hybrid approach involving write-through to
796 more persistent storage to help alleviate the issue.

797 The advantages of in-memory approaches include faster processing of intensive analysis and reporting
798 workloads. In-memory systems are especially good for analysis of real time data such as that needed for
799 some complex event processing of streams. For reporting workloads, performance improvements can
800 often be on the order of several hundred times faster—especially for sparse matrix and simulation type
801 analytics.

802 **4.4.2.2 FILE SYSTEMS**

803 Many Big Data processing frameworks and applications access their data directly from underlying file
804 systems. In almost all cases, the file systems implement some level of the Portable Operating System
805 Interface (POSIX) standards for permissions and the associated file operations. This allows other higher-
806 level frameworks for indexing or processing to operate with relative transparency as to whether the
807 underlying file system is local or fully distributed. File-based approaches consist of two layers, the file
808 system organization and the data organization within the files.

809 **4.4.2.2.1 File System Organization**

810 File systems tend to be either centralized or distributed. Centralized file systems are basically
811 implementations of local file systems that are placed on a single large storage platform (e.g., SAN or
812 NAS) and accessed via some network capability. In a virtual environment, multiple physical centralized
813 file systems may be combined, split, or allocated to create multiple logical file systems.

814 Distributed file systems (also known as cluster file systems) seek to overcome the throughput issues
815 presented by the volume and velocity characteristics of big data combine I/O throughput across multiple
816 devices (spindles) on each node, with redundancy and failover mirroring or replicating data at the block
817 level across multiple nodes. The data replication is specifically designed to allow the use of
818 heterogeneous commodity hardware across the Big Data cluster. Thus, if a single drive or an entire node
819 should fail, no data is lost because it is replicated on other nodes and throughput is only minimally
820 affected because that processing can be moved to the other nodes. In addition, replication allows for high
821 levels of concurrency for reading data and for initial writes. Updates and transaction style changes tend to
822 be an issue for many distributed file systems because latency in creating replicated blocks will create
823 consistency issues (e.g., a block is changed but another node reads the old data before it is replicated.)
824 Several file system implementations also support data compression and encryption at various levels. One
825 major caveat is that, for distributed block based file systems, the compression/encryption must be
826 splittable and allow any given block to be decompressed/ decrypted out of sequence and without access to
827 the other blocks.

828 Distributed object stores (DOSs) (also known as global object stores) are a unique example of distributed
829 file system organization. Unlike the approaches described above, which implement a traditional file
830 system hierarchy namespace approach, DOSs present a flat name space with a globally unique identifier
831 (GUID) for any given chunk of data. Generally, data in the store is located through a query against a
832 metadata catalog that returns the associated GUIDs. The GUID generally provides the underlying
833 software implementation with the storage location of the data of interest. These object stores are
834 developed and marketed for storage of very large data objects, from complete data sets to large individual
835 objects (e.g., high resolution images in the tens of gigabytes [GBs] size range). The biggest limitation of
836 these stores for Big Data tends to be network throughput (i.e., speed) because many require the object to
837 be accessed in total. However, future trends point to the concept of being able to send the
838 computation/application to the data versus needing to bring the data to the application.

839 From a maturity perspective, two key areas where distributed file systems are likely to improve are (1)
840 random write I/O performance and consistency, and (2) the generation of de facto standards at a similar or
841 greater level as the Internet Engineering Task Force (IETF) Requests for Comments (RFC), such as those
842 currently available for the network file system (NFS) protocol. DOSs, while currently available and
843 operational from several commercial providers and part of the roadmap for large organizations such as the
844 National Geospatial Intelligence Agency (NGA), currently are essentially proprietary implementations.

845 For DOSs to become prevalent within Big Data ecosystems, there should be: some level of
 846 interoperability available (i.e., through standardized APIs); standards-based approaches for data
 847 discovery; and, most importantly, standards-based approaches that allow the application to be transferred
 848 over the grid and run locally to the data versus transferring the data to the application.

849 **4.4.2.2 In File Data Organization**

850 Very little is different in Big Data for in file data organization. File based data can be text, binary data,
 851 fixed length records, or some sort of delimited structure (e.g., comma separated values [CSV], Extensible
 852 Markup Language [XML]). For record oriented storage (either delimited or fixed length), this generally is
 853 not an issue for Big Data unless individual records can exceed a block size. Some distributed file system
 854 implementations provide compression at the volume or directory level and implement it below the logical
 855 block level (e.g., when a block is read from the file system, it is decompressed/decrypted before being
 856 returned). Because of their simplicity, familiarity, and portability, delimited files are frequently the default
 857 storage format in many Big Data implementations. The tradeoff is I/O efficiency (i.e., speed). While
 858 individual blocks in a distributed file system might be accessed in parallel, each block still needs to be
 859 read in sequence. In the case of a delimited file, if only the last field of certain records is of interest with
 860 perhaps hundreds of fields, a lot of I/O and processing bandwidth is wasted.

861 Binary formats tend to be application or implementation specific. While they can offer much more
 862 efficient access due to smaller data sizes (i.e., integers are two to four bytes in binary while they are one
 863 byte per digit in ASCII [American Standard Code for Information Interchange]), they offer limited
 864 portability between different implementations. At least one popular distributed file system provides its
 865 own standard binary format, which allows data to be portable between multiple applications without
 866 additional software. However, the bulk of the indexed data organization approaches discussed below
 867 leverage binary formats for efficiency.

868 **4.4.2.3 INDEXED STORAGE ORGANIZATION**

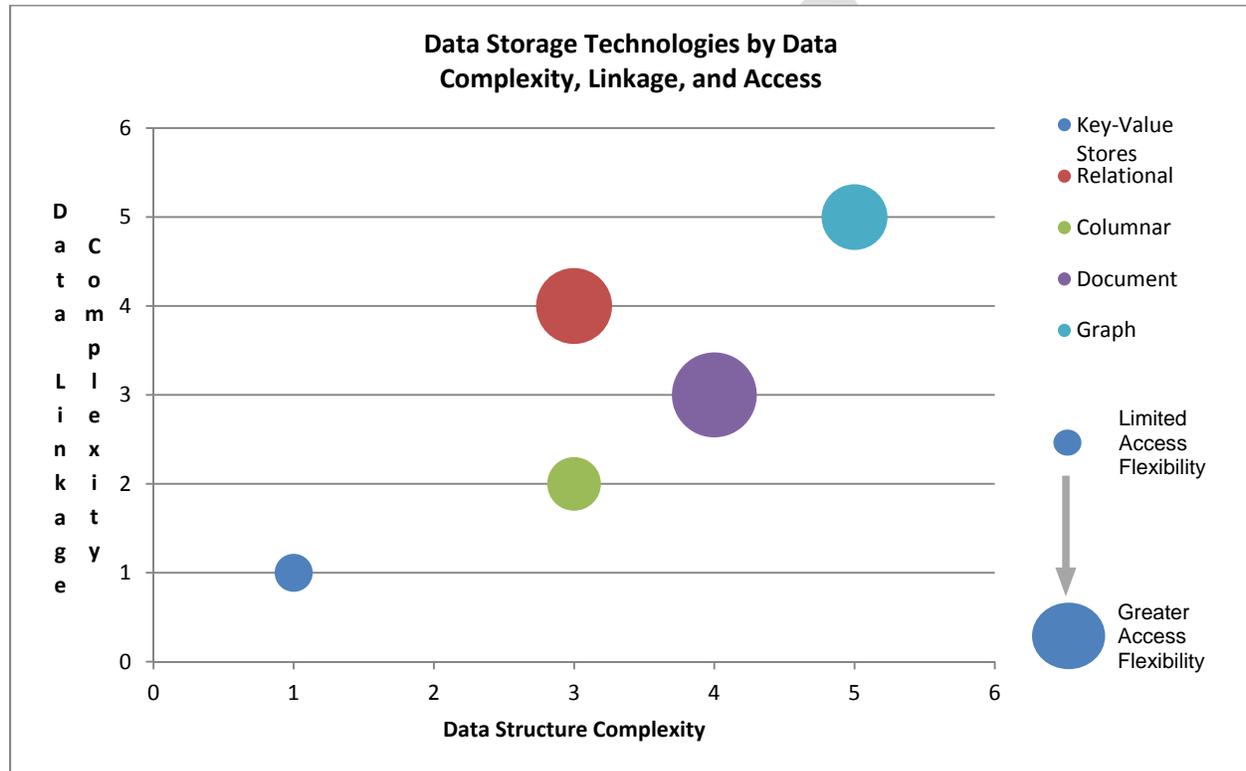
869 The very nature of Big Data (primarily the volume and velocity characteristics) practically drives
 870 requirements to some form of indexing structure. Big Data volume requires that specific data elements be
 871 located quickly without scanning across the entire dataset. Big Data velocity also requires that data can be
 872 located quickly either for matching (e.g., incoming data matches something in an existing data set) or to
 873 know where to write/update new data.

874 The choice of a particular indexing method or methods depends mostly on the data and the nature of the
 875 application to be implemented. For example, graph data (i.e., vertices, edges, and properties) can easily be
 876 represented in flat text files as vertex-edge pairs, edge-vertex-vertex triples, or vertex-edge list records.
 877 However, processing this data efficiently would require potentially loading the entire data set into
 878 memory or being able to distribute the application and data set across multiple nodes so a portion of the
 879 graph is in memory on each node. Splitting the graph across nodes requires the nodes to communicate
 880 when graph sections have vertices that connect with vertices on other processing nodes. This is perfectly
 881 acceptable for some graph applications—such as shortest path—especially when the graph is static. Some
 882 graph processing frameworks operate using this exact model. However, this approach is infeasible for
 883 large scale graphs requiring a specialized graph storage framework, where the graph is dynamic or
 884 searching or matching to a portion of the graph is needed quickly.

885 Indexing approaches tend to be classified by the features provided in the implementation, specifically: the
 886 complexity of the data structures that can be stored; how well they can process links between data; and,
 887 how easily they support multiple access patterns as shown in Figure 4. Since any of these features can be
 888 implemented in custom application code, the values portrayed represent approximate norms. For example,
 889 key-value (KV) stores work well for data that is only accessed through a single key, whose values can be
 890 expressed in a single flat structure, and where multiple records do not need to be related. While document
 891 stores can support very complex structures of arbitrary width and tend to be indexed for access via
 892 multiple document properties, they do not tend to support inter-record relationships well.

893 It is noted that the specific implementations for each storage approach vary significantly enough that all
 894 of the values for the features represented here are really ranges. For example, relational data storage
 895 implementations are supporting increasingly complex data structures and on going work aims to add more
 896 flexible access patterns natively in BigTable columnar implementations. Within Big Data, the
 897 performance of each of these features tends to drive the scalability of that approach depending on the
 898 problem being solved. For example, if the problem is to locate a single piece of data for a unique key,
 899 then KV stores will scale really well. However, if a problem requires general navigation of the
 900 relationships between multiple data records, a graph storage model will likely provide the best
 901 performance.

902



903

904

Figure 4: Data Storage Technologies

905

906 4.4.3 PROCESSING FRAMEWORKS

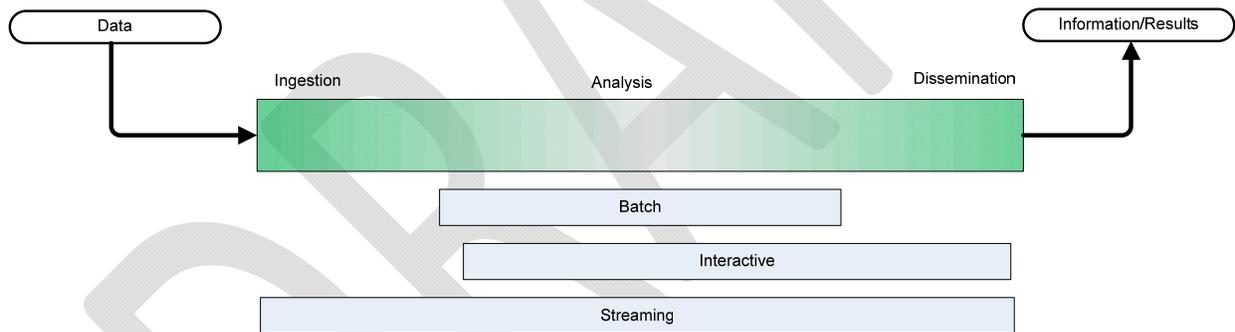
907 The processing frameworks for Big Data provide the necessary infrastructure software to support
 908 implementation of applications that can deal with the volume, velocity, variety, and variability of data.
 909 Processing frameworks define how the computation and processing of the data is organized. Big Data
 910 applications rely on various platforms and technologies to meet the challenges of scalable data analytics
 911 and operation.

912 Processing frameworks generally focus on data manipulation, which falls along a continuum between
 913 batch and streaming oriented processing. However, depending on the specific data organization platform,
 914 and actual processing requested, any given framework may support a range of data manipulation from
 915 high latency to NRT processing. Overall, many Big Data architectures will include multiple frameworks
 916 to support a wide range of requirements.

917 Typically, processing frameworks are categorized based on whether they support batch or streaming
 918 processing. This categorization is generally stated from the user perspective (e.g., how fast does a user get
 919 a response to a request). However, Big Data processing frameworks actually have three processing
 920 phases: data ingestion, data analysis, and data dissemination, which closely follow the flow of data
 921 through the architecture. The Big Data Application Provider activities control the application of specific
 922 framework capabilities to these processing phases. The batch-streaming continuum, illustrated in the
 923 processing subcomponent in Figure 2, can be applied to the three distinct processing phases. For example,
 924 data may enter a Big Data system at high velocity and the end user must quickly retrieve a summary of
 925 the prior day's data. In this case, the ingestion of the data into the system needs to be near real time
 926 (NRT) and keep up with the data stream. The analysis portion could be incremental (e.g., performed as
 927 the data is ingested) or could be a batch process performed at a specified time, while retrieval (i.e., read
 928 visualization) of the data could be interactive. Specific to the use case, data transformation may take place
 929 at any point during its transit through the system. For example, the ingestion phase may only write the
 930 data as quickly as possible, or it may run some foundational analysis to track incrementally computed
 931 information such as minimum, maximum, average. The core processing job may only perform the
 932 analytic elements required by the Big Data Application Provider and compute a matrix of data or may
 933 actually generate some rendering like a heat map to support the visualization component. To permit rapid
 934 display, the data dissemination phase almost certainly does some rendering, but the extent depends on the
 935 nature of the data and the visualization.

936 For the purposes of this discussion, most processing frameworks can be described with respect to their
 937 primary location within the information flow illustrated in Figure 5.

938



939

940

Figure 5: Information Flow

941 The green shading in Figure 5 illustrates the general sensitivity of that processing phase to latency, which
 942 is defined as the time from when a request or piece of data arrives at a system until its processing/delivery
 943 is complete. For Big Data, the ingestion may or may not require NRT performance to keep up with the
 944 data flow. Some types of analytics (specifically those categorized as complex event processing [CEP])
 945 may or may not require NRT processing. The Data Consumer generally is located at the far right of
 946 Figure 5, depending upon the use case and application batch responses (e.g., a nightly report is emailed)
 947 may be sufficient. In other cases, the user may be willing to wait minutes for the results of a query to be
 948 returned, or they may need immediate alerting when critical information arrives at the system. In general,
 949 batch analytics tend to better support long term strategic decision making, where the overall view or
 950 direction is not affected by the latest small changes in the underlying data. Streaming analytics are better
 951 suited for tactical decision making, where new data needs to be acted upon immediately. A primary use
 952 case for streaming analytics would be electronic trading on stock exchanges where the window to act on a
 953 given piece of data can be measured in microseconds. Messaging and communication provides the
 954 transfer of data between processing elements and the buffering necessary to deal with the deltas in data
 955 rate, processing times, and data requests.

956 Typically, Big Data discussions focus around the categories of batch and streaming frameworks for
 957 analytics. However, frameworks for retrieval of data that provide interactive access to Big Data are
 958 becoming a more prevalent. It is noted that the lines between these categories are not solid or distinct,
 959 with some frameworks providing aspects of each category.

960 **4.4.3.1 BATCH FRAMEWORKS**

961 Batch frameworks, whose roots stem from the mainframe processing era, are some of the most prevalent
 962 and mature components of a Big Data architecture because the historically long processing times for large
 963 data volumes. Batch frameworks ideally are not tied to a particular algorithm or even algorithm type, but
 964 rather provide a programming model where multiple classes of algorithms can be implemented. Also,
 965 when discussed in terms of Big Data, these processing models are frequently distributed across multiple
 966 nodes of a cluster. They are routinely differentiated by the amount of data sharing between
 967 processes/activities within the model.

968 In 2004, a list of algorithms for simulation in the physical sciences was developed that became known as
 969 the “Seven Dwarfs”.³ The list was recently modified and extended to the 13 algorithms (Table 2), based
 970 on the following definition: “A dwarf is an algorithmic method that computes a pattern of computation
 971 and communication.”⁴

972 **Table 2: 13 Dwarfs—Algorithms for Simulation in the Physical Sciences**

Dense Linear Algebra*	Combinational Logic
Sparse Linear Algebra*	Graph Traversal
Spectral methods	Dynamic Programming
N-Body Methods	Backtrack and Branch-and-Bound
Structured Grids*	Graphical Models
Unstructured Grids*	Finite State Machines
Map/Reduce	

973 Notes:

974 * Indicates one of the original seven dwarfs. The recent list modification removed three of the original seven algorithms: Fast
 975 Fourier Transform, Particles, and Monte Carlo.

976 Many other algorithms or processing models have been defined over the years two of the best-known
 977 models in the Big Data space, Map/Reduce (M/R) and Bulk Synchronous Parallel (BSP), are described in
 978 the following subsections.

979 **4.4.3.1.1 Map/Reduce**

980 Several major Internet search providers popularized the M/R model as they worked to implement their
 981 search capabilities. In general, M/R programs follow five basic stages:

- 982 1. Input preparation and assignment to mappers
- 983 2. Map a set of keys and values to new keys and values: $\text{Map}(k_1, v_1) \rightarrow \text{list}(k_2, v_2)$
- 984 3. Shuffle data to each reducer and each reducer sorts its input—each reducer is assigned a set of keys
 985 (k2)
- 986 4. Run the reduce on a list(v2) associated with each key and produce an output: $\text{Reduce}(k_2, \text{list}(v_2)) \rightarrow$
 987 list(v3)
- 988 5. Final output the lists(v3) from each reducer are combined and sorted by k2

989 While there is a single output, nothing in the model prohibits multiple input data sets. It is extremely
 990 common for complex analytics to be built as workflows of multiple M/R jobs. While the M/R
 991 programming model is best suited to aggregation (e.g., sum, average, group-by) type analytics, a wide
 992 variety of analytic algorithms have been implemented within processing frameworks. M/R does not

993 generally perform well with applications or algorithms that need to directly update the underlying data.
 994 For example, updating the values for a single key would require the entire data set be read, output, and
 995 then moved or copied over the original data set. Because the mappers and reducers are stateless in nature,
 996 applications that require iterative computation on parts of the data or repeated access to parts of the data
 997 set do not tend to scale or perform well under M/R.

998 Due to its shared nothing approach, the usability of M/R for Big Data applications has made it popular
 999 enough that a number of large data storage solutions (mostly those of the NoSQL variety) provide
 1000 implementations within their architecture. One major criticism of M/R early on was that the interfaces to
 1001 most implementations were at too low of a level (written in Java or JavaScript.) However many of the
 1002 more prevalent implementations now support high level procedural and declarative language interfaces
 1003 and even visual programming environments are beginning to appear.

1004 **4.4.3.1.2 Bulk Synchronous Parallel**

1005 The BSP programming model, originally developed by Leslie Valiant⁵, combines parallel processing with
 1006 the ability of processing modules to send messages to other processing modules and explicit
 1007 synchronization of the steps. A BSP algorithm is composed of what are termed “supersteps”, which are
 1008 comprised of the following three distinct elements.

- 1009 • Bulk Parallel Computation: Each processor performs the calculation/analysis on its local chunk
 1010 of data.
- 1011 • Message Passing: As each processor performs its calculations it may generate messages to other
 1012 processors. These messages are frequently updates to values associated with the local data of
 1013 other processors but may also result in the creation of additional data.
- 1014 • Synchronization: Once a processor has completed processing its local data it pauses until all
 1015 other processors have also completed their processing.

1016 This cycle can be terminated by all the processors “voting to stop”, which will generally happen when a
 1017 processor has generated no messages to other processors (e.g., no updates). All processors voting to stop
 1018 in turn indicates that there are no new updates to any processors’ data and the computation is complete.
 1019 Alternatively, the cycle may be terminated after a fixed number of supersteps have been completed (e.g.,
 1020 after a certain number of iterations of a Monte Carlo simulation).

1021 The advantage of BSP over Map/Reduce is that processing can actually create updates to the data being
 1022 processed. It is this distinction that has made BSP popular for graph processing and simulations where
 1023 computations on one node/element of data directly affect values or connections with other
 1024 nodes/elements. The disadvantage of BSP is the high cost of the synchronization barrier between
 1025 supersteps. Should the distribution of data or processing between processors become highly unbalanced
 1026 then some processors may become overloaded while others remain idle.

1027 Numerous extensions and enhancements to the basic BSP model have been developed and implemented
 1028 over the years, many of which are designed to address the balancing and cost of synchronization
 1029 problems.

1030 **4.4.3.2 STREAMING FRAMEWORKS**

1031 Streaming frameworks are built to deal with data that requires processing as fast or faster than the
 1032 velocity at which it arrives into the Big Data system. The primary goal of streaming frameworks is to
 1033 reduce the latency between the arrival of data into the system and the creation, storage, or presentation of
 1034 the results. Complex Event Processing (CEP) is one of the problem domains frequently addressed by
 1035 streaming frameworks. CEP uses data from one or more streams/sources to infer or identify events or
 1036 patterns in NRT.

1037 Almost all streaming frameworks for Big Data available today implement some form of basic workflow
 1038 processing for the streams. These workflows use messaging/communications frameworks to pass data

1039 objects (often referred to as events) between steps in the workflow. This frequently takes the form of a
 1040 directed execution graph. The distinguishing characteristics of streaming frameworks are typically
 1041 organized around the following three characteristics: event ordering and processing guarantees, state
 1042 management, and partitioning/parallelism. These three characteristics are described below.

1043 **4.4.3.2.1 Event Ordering and Processing Guarantees**

1044 This characteristic refers to whether stream processing elements are guaranteed to see messages or events
 1045 in the order they are received by the Big Data System, as well as how often a message or event may or
 1046 may not be processed. In a non-distributed and single stream mode, this type of guarantee is relatively
 1047 trivial. Once distributed and/or multiple streams are added to the system, the guarantee becomes more
 1048 complicated. With distributed processing, the guarantees must be enforced for each partition of the data
 1049 (partitioning and parallelism as further described below.) Complications arise when the process/task/job
 1050 dealing with a partition dies. Processing guarantees are typically divided into the following three classes:

- 1051 • At-most-once delivery: This is the simplest form of guarantee and allows for messages or events
 1052 to be dropped if there is a failure in processing or communications or if they arrive out of order.
 1053 This class of guarantee is applicable for data where there is no dependence of new events on the
 1054 state of the data created by prior events.
- 1055 • At-least-once delivery: Within this class, the frameworks will track each message or event (and
 1056 any downstream messages or events generated) to verify that it is processed within a configured
 1057 time frame. Messages or events that are not processed in the time allowed are re-introduced into
 1058 the stream. This mode requires extensive state management by the framework (and sometimes
 1059 the associated application) to track which events have been processed by which stages of the
 1060 workflow. However, under this class, messages or events may be processed more than once and
 1061 also may arrive out of order. This class of guarantee is appropriate for systems where every
 1062 message or event must be processed regardless of the order (e.g., no dependence on prior
 1063 events), and the application either is not affected by duplicate processing of events or has the
 1064 ability to de-duplicate events itself.
- 1065 • Exactly-once delivery: This class of framework processing requires the same top level state
 1066 tracking as At-least-once delivery but embeds mechanisms within the framework to detect and
 1067 ignore duplicates. This class often guarantees ordering of event arrivals and is required for
 1068 applications where the processing of any given event is dependent on the processing of prior
 1069 events. It is noted that these guarantees only apply to data handling within the framework. If data
 1070 is passed outside the framework processing topology, then by an application then the application
 1071 must ensure the processing state is maintained by the topology or duplicate data may be
 1072 forwarded to non-framework elements of the application.

1073 In the latter two classes, some form of unique key must be associated with each message or event to
 1074 support de-duplication and event ordering. Often, this key will contain some form of timestamp plus
 1075 the stream ID to uniquely identify each message in the stream.

1076 **4.4.3.2.2 State Management**

1077 A critical characteristic of stream processing frameworks is their ability to recover and not lose critical
 1078 data in the event of a process or node failure within the framework. Frameworks typically provide this
 1079 state management through persistence of the data to some form of storage. This persistence can be: local,
 1080 allowing the failed process to be restarted on the same node; a remote or distributed data store, allowing
 1081 the process to be restarted on any node; or, local storage that is replicated to other nodes. The tradeoff
 1082 between these storage methods is the latency introduced by the persistence. Both the amount of state data
 1083 persisted and the time required to assure that the data is persisted contribute to the latency. In the case of a
 1084 remote or distributed data store, the latency required is generally dependent on the extent to which the
 1085 data store implements ACID (Atomicity, Consistency, Isolation, Durability) or BASE (Basically
 1086 Available, Soft state, Eventual consistency) style consistency. With replication of local storage, the

1087 reliability of the state management is entirely tied to the ability of the replication to recover in the event of
 1088 a process or node failure. Sometimes this state replication is actually implemented using the same
 1089 messaging/communication framework that is used to communicate with and between stream processors.
 1090 Some frameworks actually support full transaction semantics, including multi-stage commits and
 1091 transaction rollbacks. The tradeoff is the same one that exists for any transaction system is that any type
 1092 of ACID-like guarantee will introduce latency. Too much latency at any point in the stream flow can
 1093 create bottlenecks and, depending on the ordering or processing guarantees, can result in deadlock or loop
 1094 states—especially when some level of failure is present.

1095 **4.4.3.2.3 Partitioning and Parallelism**

1096 This streaming framework characteristic relates to the distribution of data across nodes and worker tasks
 1097 to provide the horizontal scalability needed to address the volume and velocity of Big Data streams. This
 1098 partitioning scheme must interact with the resource management framework to allocate resources. The
 1099 even distribution of data across partitions is essential so that the associated work is evenly distributed.
 1100 The even data distribution directly relates to selection of a key (e.g., user ID, host name) that can be
 1101 evenly distributed. The simplest form might be using a number that increments by one and then is
 1102 processed with a modulus function of the number of tasks/workers available. If data dependencies require
 1103 all records with a common key be processed by the same worker, then assuring an even data distribution
 1104 over the life of the stream can be difficult. Some streaming frameworks address this issue by supporting
 1105 dynamic partitioning where the partition of overloaded workers is split and allocated to existing workers
 1106 or newly created workers. To achieve success—especially with a data/state dependency related to the
 1107 key—it is critical that the framework have state management, which allows the associated state data to be
 1108 moved/transitioned to the new/different worker.

1109 **4.4.4 MESSAGING/COMMUNICATIONS FRAMEWORKS**

1110 Messaging and communications frameworks have their roots in the High Performance Computing (HPC)
 1111 environments long popular in the scientific and research communities. Messaging/Communications
 1112 Frameworks were developed to provide APIs for the reliable queuing, transmission, and receipt of data
 1113 between nodes in a horizontally scaled cluster. These frameworks typically implement either a point-to-
 1114 point transfer model or a store-and-forward model in their architecture. Under a point-to-point model,
 1115 data is transferred directly from the sender to the receivers. The majority of point-to-point
 1116 implementations do not provide for any form of message recovery should there be a program crash or
 1117 interruption in the communications link between sender and receiver. These frameworks typically
 1118 implement all logic within the sender and receiver program space, including any delivery guarantees or
 1119 message retransmission capabilities. One common variation of this model is the implementation of
 1120 multicast (i.e., one-to-many or many-to-many distribution), which allows the sender to broadcast the
 1121 messages over a “channel”, and receivers in turn listen to those channels of interest. Typically, multicast
 1122 messaging does not implement any form of guaranteed receipt. With the store-and-forward model, the
 1123 sender would address the message to one or more receivers and send it to an intermediate broker, which
 1124 would store the message and then forward it on to the receivers. Many of these implementations support
 1125 some form of persistence for messages not yet delivered, providing for recovery in the event of process or
 1126 system failure. Multicast messaging can also be implemented in this model and is frequently referred to as
 1127 a pub/sub model.

1128 **4.4.5 RESOURCE MANAGEMENT FRAMEWORK**

1129 As Big Data systems have evolved and become more complex, and as businesses work to leverage limited
 1130 computation and storage resources to address a broader range of applications and business challenges, the
 1131 requirement to effectively manage those resources has grown significantly. While tools for resource
 1132 management and “elastic computing” have expanded and matured in response to the needs of cloud
 1133 providers and virtualization technologies, Big Data introduces unique requirements for these tools.

1134 However, Big Data frameworks tend to fall more into a distributed computing paradigm, which presents
1135 additional challenges.

1136 The Big Data characteristics of volume and velocity drive the requirements with respect to Big Data
1137 resource management. Elastic computing (i.e., spawning another instance of some service) is the most
1138 common approach to address expansion in volume or velocity of data entering the system. CPU and
1139 memory are the two resources that tend to be most essential to managing Big Data situations. While
1140 shortages or over-allocation of either will have significant impacts on system performance, improper or
1141 inefficient memory management is frequently catastrophic. Big Data differs and becomes more complex
1142 in the allocation of computing resources to different storage or processing frameworks that are optimized
1143 for specific applications and data structures. As such, resource management frameworks will often use
1144 data locality as one of the input variables in determining where new processing framework elements (e.g.,
1145 master nodes, processing nodes, job slots) are instantiated. Importantly, because the data is big (i.e., large
1146 volume), it generally is not feasible to move data to the processing frameworks. In addition, while nearly
1147 all Big Data processing frameworks can be run in virtualized environments, most are designed to run on
1148 bare metal commodity hardware to provide efficient I/O for the volume of the data.

1149 Two distinct approaches to resource management in Big Data frameworks are evolving. The first is intra-
1150 framework resource management, where the framework itself manages allocation of resources between its
1151 various components. This allocation is typically driven by the framework's workload and often seeks to
1152 "turn off" unneeded resources to either minimize overall demands of the framework on the system or to
1153 minimize the operating cost of the system by reducing energy use. With this approach, applications can
1154 seek to schedule and request resources that—much like main frame operating systems of the past—are
1155 managed through scheduling queues and job classes.

1156 The second approach is inter-framework resource management, which is designed to address the needs of
1157 many Big Data systems to support multiple storage and processing frameworks that can address and be
1158 optimized for a wide range of applications. With this approach, the resource management framework
1159 actually runs as a service that supports and manages resource requests from frameworks, monitoring
1160 framework resource usage, and in some cases manages application queues. In many ways, this approach
1161 is like the resource management layers common in cloud/virtualization environments, and there are
1162 efforts underway to create hybrid resource management frameworks that handle both physical and virtual
1163 resources.

1164 Taking these concepts further and combining them is resulting in the emerging technologies built around
1165 what is being termed software-defined data centers (SDDCs). This expansion on elastic and cloud
1166 computing goes beyond the management of fixed pools of physical resources as virtual resources to
1167 include the automated deployment and provisioning of features and capabilities onto physical resources.
1168 For example, automated deployment tools that interface with virtualization or other framework APIs can
1169 be used to automatically stand up entire clusters or to add additional physical resources to physical or
1170 virtual clusters.

1171 **4.5 DATA CONSUMER**

1172 Similar to the Data Provider, the role of Data Consumer within the NBDRA can be an actual end user or
1173 another system. In many ways, this role is the mirror image of the Data Provider, with the entire Big Data
1174 framework appearing like a Data Provider to the Data Consumer. The activities associated with the Data
1175 Consumer role include the following:

- 1176 • Search and Retrieve
- 1177 • Download
- 1178 • Analyze Locally
- 1179 • Reporting

- 1180 • Visualization
- 1181 • Data to Use for Their Own Processes

1182 The Data Consumer uses the interfaces or services provided by the Big Data Application Provider to get
1183 access to the information of interest. These interfaces can include data reporting, data retrieval, and data
1184 rendering.

1185 This role will generally interact with the Big Data Application Provider through its access function to
1186 execute the analytics and visualizations implemented by the Big Data Application Provider. This
1187 interaction may be demand-based, where the Data Consumer initiates the command/transaction and the
1188 Big Data Application Provider replies with the answer. The interaction could include interactive
1189 visualizations, creating reports, or drilling down through data using business intelligence functions
1190 provided by the Big Data Application Provider. Alternately, the interaction may be stream- or push-based,
1191 where the Data Consumer simply subscribes or listens for one or more automated outputs from the
1192 application. In almost all cases, the Security and Privacy fabric around the Big Data architecture would
1193 support the authentication and authorization between the Data Consumer and the architecture, with either
1194 side able to perform the role of authenticator/authorizer and the other side providing the credentials. Like
1195 the interface between the Big Data architecture and the Data Provider, the interface between the Data
1196 Consumer and Big Data Application Provider would also pass through the three distinct phases of
1197 initiation, data transfer, and termination.

1198

1199 **5 MANAGEMENT FABRIC OF THE NBDRA**

1200 The Big Data characteristics of volume, velocity, variety, and variability demand a versatile management
 1201 platform for storing, processing, and managing complex data Management of Big Data systems should
 1202 handle both system and data related aspects of the Big Data environment. The Management Fabric of the
 1203 NBDRA encompasses two general groups of activities: system management and Big Data lifecycle
 1204 management. System management includes activities such as provisioning, configuration, package
 1205 management, software management, backup management, capability management, resources
 1206 management, and performance management. Big Data lifecycle management involves activities
 1207 surrounding the data lifecycle of collection, preparation/curation, analytics, visualization, and access.

1208 As discussed above, the NBDRA represents a broad range of Big Data systems—from tightly-coupled
 1209 enterprise solutions integrated by standard or proprietary interfaces to loosely-coupled vertical systems
 1210 maintained by a variety of stakeholders or authorities bound by agreements, standard interfaces, or de
 1211 facto standard interfaces. Therefore, different considerations and technical solutions would be applicable
 1212 for different cases.

1213 **5.1 SYSTEM MANAGEMENT**

1214 The characteristics of Big Data pose system management challenges on traditional management
 1215 platforms. To efficiently capture, store, process, analyze, and distribute complex and large datasets
 1216 arriving or leaving with high velocity, a resilient system management is needed.

1217 As in traditional systems, system management for Big Data architecture involves provisioning,
 1218 configuration, package management, software management, backup management, capability
 1219 management, resources management, and performance management of the Big Data infrastructure,
 1220 including compute nodes, storage nodes, and network devices. Due to the distributed and complex nature
 1221 of the Big Data infrastructure, system management for Big Data is challenging, especially with respect to
 1222 the capability for controlling, scheduling, and managing the processing frameworks to perform the
 1223 scalable, robust, and secure analytics processing required by the Big Data Application Provider. The Big
 1224 Data infrastructure may contain SAN or NAS storage devices, cloud storage spaces, NoSQL databases,
 1225 M/R clusters, data analytics functions, search and indexing engines, and messaging platforms. The
 1226 supporting enterprise computing infrastructure can range from traditional data centers, cloud services, and
 1227 dispersed computing nodes of a grid. To manage the distributed and complex nature of the Big Data
 1228 infrastructure, system management relies on the following:

- 1229 • Standard protocols such as SNMP which are used to transmit status about resources and fault
 1230 information to the management fabric components
- 1231 • Deployable agents or management connectors which allow the management fabric to both
 1232 monitor and also control elements of the framework.

1233 These two items aid in monitoring the health of various types of computing resources and coping with
 1234 performance and failures incidents while maintaining the quality of service (QoS) levels required by the
 1235 Big Data Application Provider. Management connectors are necessary for scenarios where the cloud
 1236 service providers expose management capabilities via APIs. It is conceivable that the infrastructure
 1237 elements contain autonomic, self-tuning, and self-healing capabilities, thereby reducing the centralized
 1238 model of system management. In large infrastructures with many thousands of computing and storage
 1239 nodes, the provisioning of tools and applications should be as automated as possible. Software
 1240 installation, application configuration, and regular patch maintenance should be pushed out and replicated
 1241 across the nodes in an automated fashion, which could be done based on the topology knowledge of the

1242 infrastructure. With the advent of virtualization, the utilization of virtual images may speed up the
1243 recovery process and provide efficient patching that can minimize downtime for scheduled maintenance.

1244 In an enterprise environment, the management platform would typically provide enterprise-wide
1245 monitoring and administration of the Big Data distributed components. This includes network
1246 management, fault management, configuration management, system accounting, performance
1247 management, and security management.

1248 In a loosely-coupled vertical system, each independent stakeholder is responsible for its own system
1249 management, security, and integration. Each stakeholder is responsible for integration within the Big Data
1250 distributed system using the interfaces provided by other stakeholders.

1251 5.2 BIG DATA LIFECYCLE MANAGEMENT

1252 Big Data lifecycle management (BDLM) faces more challenges compared to traditional data lifecycle
1253 management (DLM), which may require less data transfer, processing, and storage. However, BDLM still
1254 inherits the DLM phases in terms of data acquisition, distribution, use, migration, maintenance, and
1255 disposition—but at a much bigger processing scale. Big Data Application Providers may require much
1256 more computational processing for collection, preparation/curation, analytics, visualization, and access to
1257 be able to use the analytic results. In other words, the BDLM activity includes verification that the data
1258 are handled correctly by other NBDRA components in each process within the data lifecycle—from the
1259 moment they are ingested into the system by the Data Provider, until the data are processed or removed
1260 from the system.

1261 The importance of BDLM to Big Data is demonstrated through the following considerations:

- 1262 • Data volume can be extremely large that may overwhelm the storage capacity, or make storing
1263 incoming data prohibitively expensive
- 1264 • Data velocity, the rate at which data can be captured and ingested into the system, can
1265 overwhelm available storage space at a given time. Even with the elastic storage service
1266 provided by cloud computing for handling dynamic storage needs, uncontrolled data
1267 management may also be unnecessarily costly for certain application requirements.
- 1268 • Different Big Data applications will likely have different requirements for the lifetime of a piece
1269 of data. The differing requirements have implications on how often data must be refreshed so
1270 that processing results are valid and useful. In data refreshment, old data are dispositioned and
1271 not fed into analytics or discovery programs. At the same time, new data is ingested and taken
1272 into account by the computations. For example, real-time applications will need very short data
1273 lifetime but a market study of consumers' interest in a product line may need to mine data
1274 collected over a longer period of time

1275 Because the task of BDLM can be distributed among different organizations and/or individuals within the
1276 Big Data computing environment, coordination of data processing between NBDRA components has
1277 greater difficulty in complying with policies, regulations, and security requirements. Within this context,
1278 BDLM may need to include the following subactivities:

- 1279 • **Policy Management:** Captures the requirements for the data lifecycle that allows old data to be
1280 dispositioned and new data to be considered by Big Data applications; maintains the migration
1281 and disposition strategies that specify the mechanism for data transformation and dispositioning,
1282 including transcoding data, transferring old data to lower tier storage for archival purpose,
1283 removing data, or marking data as in situ
- 1284 • **Metadata Management:** Enables BDLM, since metadata are used to store information that
1285 governs the management of the data within the system. Essential metadata information includes
1286 persistent identification of the data, fixity/quality, and access rights. The challenge is to find the
1287 minimum set of elements to execute the required BDLM strategy in an efficient manner.

- 1288
- 1289
- 1290
- 1291
- 1292
- 1293
- 1294
- 1295
- 1296
- 1297
- 1298
- 1299
- 1300
- 1301
- 1302
- 1303
- 1304
- 1305
- 1306
- 1307
- 1308
- 1309
- **Accessibility Management** Change of data accessibility over time: For example, census data can be made available to the public after 72 years.⁶ BDLM is responsible for triggering the accessibility update of the data or sets of data according to policy and legal requirements. Normally, data accessibility information is stored in the metadata.
 - **Data Recovery:** BDLM can include the recovery of data that were lost due to disaster or system/storage fault. Traditionally, data recovery can be achieved using regular backup and restore mechanisms. However, given the large volume of Big Data, traditional backup may not be feasible. Instead, replication may have to be designed within the Big Data ecosystem. Depending on the tolerance of data loss—each application has its own tolerance level—replication strategies have to be designed. The replication strategy includes the replication window time, the selected data to be replicated, and the requirements for geographic disparity. Additionally, in order to cope with the large volume of Big Data, data backup and recovery should consider the use of modern technologies within the Big Data Framework Provider.
 - **Preservation Management:** The system maintains data integrity so that the veracity and velocity of the analytics process are fulfilled. Due to the extremely large volume of Big Data, preservation management is responsible for disposition-aged data contained in the system. Depending on the retention policy, these aged data can be deleted or migrated to archival storage. In the case where data must be retained for years, decades, and even centuries, a preservation strategy will be needed so the data can be accessed by the provider components if required. This will invoke long-term digital preservation that can be performed by Big Data Application Providers using the resources of the Big Data Framework Provider.

1310 In the context of Big Data, BDLM contends with the Big Data characteristics of volume, velocity, variety,
1311 and variability. As such, BDLM and its subactivities interact with other components of the NBDRA as
1312 shown in the following examples:

- 1313
- 1314
- 1315
- 1316
- 1317
- 1318
- 1319
- 1320
- 1321
- 1322
- 1323
- 1324
- 1325
- 1326
- 1327
- 1328
- 1329
- 1330
- **System Orchestrator:** BDLM enables data scientists to initiate any combination of processing including accessibility management, data backup/recovery, and preservation management. The process may involve other components of the NBDRA, such as Big Data Application Provider and Big Data Framework Provider. For example, data scientists may want to interact with the Big Data Application Provider for data collection and curation, invoke the Big Data Framework Provider to perform certain analysis, and grant access to certain users to access the analytic results from the Data Consumer.
 - **Data Provider:** BDLM manages ingestion of data and metadata from the data source(s) into the Big Data system, which may include logging the entry event in the metadata by the Data Provider
 - **Big Data Application Provider:** BDLM executes data masking and format transformations for data preparation or curation purpose
 - **Big Data Framework Provider:** BDLM executes basic bit-level preservation and data backup and recovery according to the recovery strategy
 - **Data Consumer:** BDLM ensures that relevant data and analytic results are available with proper access control for consumers and software agents to consume within the BDLM policy strategy
 - **Security and Privacy Fabric:** Keeps the BDLM up to date according to new security policy and regulations

1331 The Security and Privacy Fabric also uses information coming from BDLM with respect to data
1332 accessibility. The Security and Privacy Fabric control access to the functions and data usage produced by
1333 the Big Data system. This data access control can be informed by the can use metadata, which is managed
1334 and updated by BDLM.

1335

1336

6 SECURITY AND PRIVACY FABRIC OF THE NBDRA

1337 Security and privacy considerations form a fundamental aspect of the NBDRA. This is geometrically
1338 depicted in Figure 2 by the Security and Privacy Fabric surrounding the five main components, indicating
1339 that all components are affected by security and privacy considerations. Thus, the role of security and
1340 privacy is correctly depicted in relation to the components but does not expand into finer details, which
1341 may be more accurate but are best relegated to a more detailed security and privacy reference
1342 architecture. The Data Provider and Data Consumer are included in the Security and Privacy Fabric since,
1343 at the least, they may often nominally agree on security protocols and mechanisms. The Security and
1344 Privacy Fabric is an approximate representation that alludes to the intricate interconnected nature and
1345 ubiquity of security and privacy throughout the NBDRA. Additional details about the Security and
1346 Privacy Fabric are included in the *NIST Interoperability Framework: Volume 4, Security and Privacy*
1347 document.

1348

DRAFT

1349 7 CONCLUSION

1350 The NBD-PWG Reference Architecture Subgroup prepared this *NIST Big Data Interoperability*
 1351 *Framework: Volume 6, Reference Architecture* to provide a vendor-neutral, technology- and
 1352 infrastructure-agnostic conceptual model and examine related issues. The conceptual model, referred to as
 1353 the NIST Big Data Reference Architecture (NBDRA), was a collaborative effort within the Subgroup and
 1354 with the other NBD-PWG subgroups. The goal of the NBD-PWG Reference Architecture Subgroup is to
 1355 develop an open reference architecture for Big Data that achieves the following objectives:

- 1356 • Provides a common language for the various stakeholders
- 1357 • Encourages adherence to common standards, specifications, and patterns
- 1358 • Provides consistent methods for implementation of technology to solve similar problem sets
- 1359 • Illustrates and improves understanding of the various Big Data components, processes, and
 1360 systems, in the context of a vendor- and technology- agnostic Big Data conceptual model
- 1361 • Provides a technical reference for U.S. government departments, agencies, and other consumers
 1362 to understand, discuss, categorize, and compare Big Data solutions
- 1363 • Facilitates analysis of candidate standards for interoperability, portability, reusability, and
 1364 extensibility

1365 This document presents the results of the first stage of NBDRA development. The *NIST Big Data*
 1366 *Interoperability Framework* will be released in three versions, which correspond to the three stages of the
 1367 NBD-PWG work. The three stages aim to achieve the following:

1368 Stage 1: Identify the common reference architecture components of Big Data implementations and
 1369 formulate the technology-independent NBDRA

1370 Stage 2: Define general interfaces between the NBDRA components

1371 Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces

1372 This document (Version 1) presents the overall NBDRA components and fabrics with high-level
 1373 description and functionalities.

1374 Version 2 activities will focus on the definition of general interfaces between the NBDRA components by
 1375 achieving the following:

- 1376 • Select use cases from the 62 (51 general and 11 security and privacy) submitted use cases or
 1377 other, to be identified, meaningful use cases
- 1378 • Work with domain experts to identify workflow and interactions from the System Orchestrator
 1379 to the rest of the NBDRA components
- 1380 • Implement small scale, manageable, and well-defined confined environment and model
 1381 interaction between NBDRA components and fabrics
- 1382 • Aggregate the common data workflow and interaction between NBDRA components/fabrics and
 1383 package them into general interfaces

1384 Version 3 activities will focus on validation of the NBDRA through the use of the defined NBDRA
 1385 interfaces to build general Big Data applications. The validation strategy will include the following:

- 1386 • Implement the same set of use cases used in Version 2 by using the defined general interfaces
- 1387 • Identify and implement a few new use cases outside the Version 2 scenarios
- 1388 • Enhance general NBDRA interfaces through lessons learned from the implementations in
 1389 Version 3 activities

1390 The general interfaces developed during Version 3 activities will offer a starting point for further
1391 refinement by any interested parties and is not intended to be a definitive solution to address all
1392 implementation needs.
1393

DRAFT

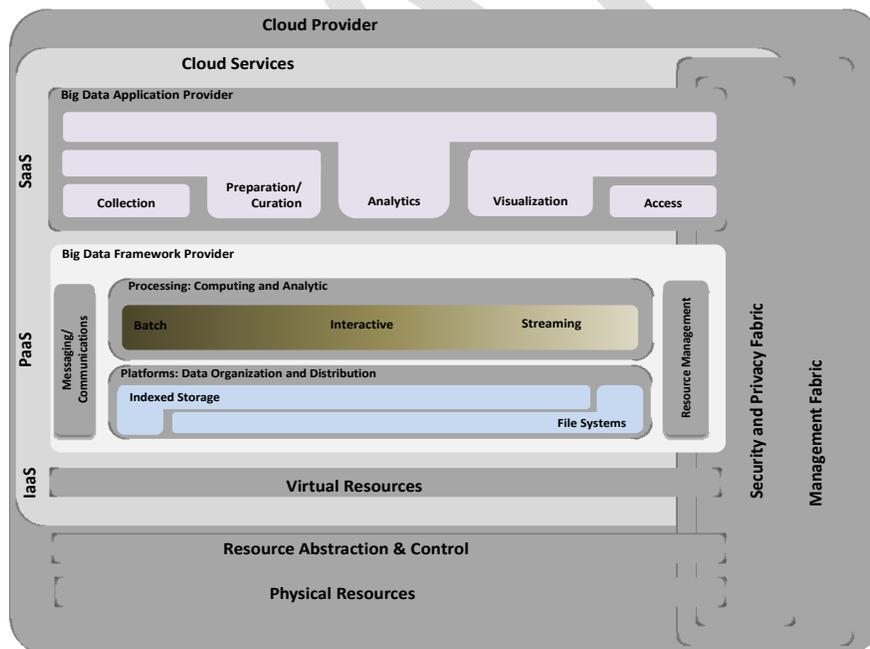
1394

Appendix A: Deployment Considerations

1395 The NIST Big Data Reference Architecture is applicable to a variety of business environments and
 1396 technologies. As a result, possible deployment models are not part of the core concepts discussed in the
 1397 main body of this document. However, the loosely coupled and distributed natures of Big Data
 1398 Framework Provider functional components allows it to be deployed using multiple infrastructure
 1399 elements as described in Section 4.4.1. The two most common deployment configurations are directly on
 1400 physical resources or on top of an IaaS cloud computing framework. The choices between these two
 1401 configurations are driven by needs of efficiency/performance and elasticity. Physical infrastructures are
 1402 typically used to obtain predictable performance and efficient utilization of CPU and I/O bandwidth since
 1403 it eliminates the overhead and additional abstraction layers typical in the virtualized environments for
 1404 most IaaS implementations. IaaS cloud based deployments on are typically used when elasticity is needed
 1405 to support changes in workload requirements. The ability to rapidly instantiate additional processing
 1406 nodes or framework components allows the deployment to adapt to either increased or decreased
 1407 workloads. By allowing the deployment footprint to grow or shrink based on workload demands this
 1408 deployment model can provide cost savings when public or shared cloud services are used and more
 1409 efficient use and energy consumption when a private cloud deployment is used. Recently, a hybrid
 1410 deployment model known as Cloud Bursting has become popular. In this model a physical deployment is
 1411 augmented by either public or private IaaS cloud services. When additional processing is needed to
 1412 support the workload additional the additional framework component instances are established on the
 1413 IaaS infrastructure and then deleted when no longer require.

1414

Figure A-1: Big Data Framework Deployment Options



1415 In addition to providing IaaS support cloud providers are now offering Big Data Frameworks under a
 1416 PaaS model. Under this model the system implementer is freed from the need to establish and manage the
 1417 complex configuration and deployment typical of many Big Data Framework components. The
 1418 implementer simply needs to specify the size of the cluster required and the cloud provider manages the
 1419 provisioning, configuration, and deployment of all the framework components. There are even some
 1420 nascent offerings for specialized SaaS Big Data applications appearing in the market that implement the
 1421 Big Data Application Provider functionality within the cloud environment. Figure A-1 illustrates how the

1422 components of the NBDRA might align onto the NIST Cloud Reference architecture⁷. The following
 1423 sections describe some of the high level interactions required between the Big Data Architecture elements
 1424 and the Cloud Service Provider elements.

1425 **CLOUD SERVICE PROVIDERS**

1426 Recent data analytics solutions use algorithms that can utilize and benefit from the frameworks of the
 1427 cloud computing systems. Cloud computing has essential characteristics such as rapid elasticity and
 1428 scalability, multi-tenancy, on-demand self-service and resource pooling, which together can significantly
 1429 lower the barriers to the realization of Big Data implementations.

1430 The **Cloud Service Provider (CSP)** implements and delivers **cloud services**. Processing of a service
 1431 invocation is done by means of an instance of the service implementation, which may involve the
 1432 composition and invocation of other services as determined by the design and configuration of the service
 1433 implementation.

1434 ***Cloud Service Component***

1435 The cloud service component contains the implementation of the cloud services provided by a CSP. It
 1436 contains and controls the software components that implement the services (but not the underlying
 1437 hypervisors, host operating systems, device drivers, etc.)

1438 Cloud services can be described in terms of service categories.

1439 Cloud services are also grouped into categories, where each service category is characterized by qualities
 1440 that are common between the services within the category. The NIST Cloud Computing Reference Model
 1441 defines the following cloud service categories:

- 1442 • Infrastructure as a Services (IaaS)
- 1443 • Platform as a Service (PaaS)
- 1444 • Software as a Service (SaaS)

1445 ***Resource Abstraction and Control Component***

1446 The Resource Abstraction and Control component is used by cloud service providers to provide access to
 1447 the physical computing resources through software abstraction. Resource abstraction needs to assure
 1448 efficient, secure, and reliable usage of the underlying physical resources. The control feature of the
 1449 component enables the management of the resource abstraction features.

1450 The Resource Abstraction and Control component enables a cloud service provider to offer qualities such
 1451 as rapid elasticity, resource pooling, on-demand self-service and scale-out. The Resource Abstraction and
 1452 Control component can include software elements such as hypervisors, virtual machines, virtual data
 1453 storage, and time-sharing.

1454 The Resource Abstraction and Control component enables control functionality. For example, there may
 1455 be a centralized algorithm to control, correlate, and connect various processing, storage, and networking
 1456 units in the physical resources so that together they deliver an environment where IaaS, PaaS or SaaS
 1457 cloud service categories can be offered. The controller might decide which CPUs/racks contain which
 1458 virtual machines executing which parts of a given cloud workload, and how such processing units are
 1459 connected to each other, and when to dynamically and transparently reassign parts of the workload to new
 1460 units as conditions change.

1461 ***Security and Privacy and Management Functions***

1462 In almost all cases the Cloud Provider will provide elements of the Security, Privacy, and Management
 1463 functions. Typically the provider will support high level security/privacy functions that control access to
 1464 the Big Data Applications and Frameworks while the frameworks themselves must control access to their

1465 underlying data and application services. Many times the Big Data specific functions for security and
1466 privacy will depend on and must interface with functions provided by the cloud service provider.
1467 Similarly, management functions are often split between the Big Data implementation and the Cloud
1468 Provider implementations. Here the cloud provider would handle the deployment and provisioning of Big
1469 Data architecture elements within its IaaS infrastructure. The cloud provider may provide high level
1470 monitoring functions to allow the Big Data implementation to track performance and resource usage of its
1471 components. In, many cases the Resource Management element of the Big Data Framework will need to
1472 interface to the Cloud Service Provider's management framework to request additional resources.

1473 **PHYSICAL RESOURCE DEPLOYMENTS**

1474 As stated above deployment on physical resources is frequently used when performance characteristics
1475 are paramount. The nature of the underlying physical resource implementations to support Big Data
1476 requirements has evolved significantly over the years. Specialized, high performance super computers
1477 with custom approaches for sharing resources (memory, CPU, storage) between nodes has given way to
1478 shared nothing computing clusters built from commodity servers. The custom super computing
1479 architectures almost always required custom development and components to take advantage of the
1480 shared resources. The commodity server approach both reduced the hardware investment and allowed the
1481 Big Data frameworks to provide higher level abstractions for the sharing and management of resources in
1482 the cluster. The Recent trends now involve density, power, cooling optimized server form factors that
1483 seek to maximize the available computing resources while minimizing size, power and/or cooling
1484 requirements. This approach retains the abstraction and portability advantages of the shared nothing
1485 approaches while providing improved efficiency.

1486

1487

1488 Appendix B: Terms and Definitions

1489 NBDRA COMPONENTS

- 1490 • **Big Data Engineering:** Advanced techniques that harness independent resources for building
1491 scalable data systems when the characteristics of the datasets require new architectures for
1492 efficient storage, manipulation, and analysis.
- 1493 • **Data Provider:** Organization or entity that introduces information feeds into the Big Data
1494 system for discovery, access, and transformation by the Big Data system.
- 1495 • **Big Data Application Provider:** Organization or entity that executes a generic vertical system
1496 data lifecycle, including: (a) data collection from various sources, (b) multiple data
1497 transformations being implemented using both traditional and new technologies, (c) diverse data
1498 usage, and (d) data archiving.
- 1499 • **Big Data Framework Provider:** Organization or entity that provides a computing fabric (such
1500 as system hardware, network, storage, virtualization, and computing platform) to execute certain
1501 Big Data applications, while maintaining security and privacy requirements
- 1502 • **Data Consumer:** End users or other systems that use the results of data applications.
- 1503 • **System Orchestrator:** Organization or entity that defines and integrates the required data
1504 transformations components into an operational vertical system.

1505 OPERATIONAL CHARACTERISTICS

- 1506 • Interoperability: The capability to communicate, to execute programs, or to transfer data among
1507 various functional units under specified conditions.
- 1508 • Portability: The ability to transfer data from one system to another without being required to
1509 recreate or reenter data descriptions or to modify significantly the application being transported.
- 1510 • Privacy: The assured, proper, and consistent collection, processing, communication, use and
1511 disposition of data associated with personal information (PI) and personally-identifiable
1512 information (PII) throughout its lifecycle.
- 1513 • Security: Protecting data, information, and systems from unauthorized access, use, disclosure,
1514 disruption, modification, or destruction in order to provide:
 - 1515 ○ Integrity: guarding against improper data modification or destruction, and includes ensuring
1516 data nonrepudiation and authenticity
 - 1517 ○ Confidentiality: preserving authorized restrictions on access and disclosure, including means
1518 for protecting personal privacy and proprietary data
 - 1519 ○ Availability: ensuring timely and reliable access to and use of data
- 1520 • Elasticity: The ability to dynamically scale up and down as a real time response to the workload
1521 demand. Elasticity will depend on the Big Data system, but adding or removing "software
1522 threads" and "virtual or physical servers" are two widely used scaling techniques. Many types of
1523 workload demands drive elastic responses, including Web based users, software agents, and
1524 periodic batch jobs.
- 1525 • Persistence: The placement/storage of data in a medium design to allow its future access.

1526 PROVISIONING MODELS

- 1527 • Infrastructure as a Service (IaaS): The capability provided to the consumer to provision
1528 processing, storage, networks, and other fundamental computing resources where the consumer
1529 is able to deploy and run arbitrary software, which can include operating systems and
1530 applications. The consumer does not manage or control the underlying cloud infrastructure but

- 1531 has control over operating systems, storage, deployed applications, and possibly limited control
1532 of select networking components (e.g., host firewalls).⁸
- 1533 • Platform as a Service (PaaS): The capability provided to the consumer to deploy onto the cloud
1534 infrastructure consumer-created or acquired applications created using programming languages
1535 and tools supported by the provider. The consumer does not manage or control the underlying
1536 cloud infrastructure including network, servers, operating systems, or storage, but has control
1537 over the deployed applications and possibly application hosting environment configurations.
1538 (Source: NIST CC Definition)
 - 1539 • Software as a Service (SaaS): The capability provided to the consumer to use applications
1540 running on a cloud infrastructure. The consumer does not manage or control the underlying
1541 cloud infrastructure including network, servers, operating systems, storage, or even individual
1542 application capabilities, with the possible exception of limited user-specific application
1543 configuration settings. (Source: NIST CC Definition)
- 1544

DRAFT

1545 **Appendix C: Examples of Big Data Organization** 1546 **Approaches**

1547 This appendix provides an overview of several common Big Data Organization Approaches. The reader
1548 should keep in mind that new and innovative approaches are emerging regularly and that some of these
1549 approaches are hybrid models that combine features of several indexing techniques (e.g., relational and
1550 columnar, or relational and graph).

1551 **RELATIONAL STORAGE MODELS**

1552 This model is perhaps the most familiar to folks as the basic concept has existed since the 1950s and the
1553 Structured Query Language (SQL) is a mature standard for manipulating (search, insert, update, delete)
1554 relational data. In the relational model, data is stored as rows with each field representing a column
1555 organized into Table based on the logical data organization. The problem with relational storage models
1556 and Big Data is the join between one or more tables. While the size of 2 or more tables of data
1557 individually might be small the join (or relational matches) between those tables will generate
1558 exponentially more records. The appeal of this model for organizations just adopting Big Data is its
1559 familiarity. The pit falls are some of the limitations and more importantly the tendency to adopt standard
1560 RDBMS practices (high normalization, detailed and specific indexes) and performance expectations

1561 Big data implementations of relational storage models are relatively mature and have been adopted by a
1562 number of organizations. They are also maturing very rapidly with new implementations focusing on
1563 improved response time. Many Big Data implementations take a brute force approach to scaling relational
1564 queries. Essentially, queries are broken into stages but more importantly processing of the input tables is
1565 distributed across multiple nodes (often as a Map/Reduce job). The actual storage of the data can be flat
1566 files (delimited or fixed length) where each record/line in the file represents a row in a table. Increasingly
1567 however these implementations are adopting binary storage formats optimized for distributed file
1568 systems. These formats will often use block level indexes and column oriented organization of the data to
1569 allow individual fields to be accessed in records without needing to read the entire record. Despite this,
1570 most Big Data Relational storage models are still “batch oriented” systems designed for very complex
1571 queries which generate very large intermediate cross-product matrices from joins so even the simplest
1572 query can required 10s of seconds to complete. There is significant work going on and emerging
1573 implementations that are seeking to provide a more interactive response and interface.

1574 Early implementations only provided limited data types and little or no support for indexes. However,
1575 most current implementations have support for complex data structures and basic indexes. However,
1576 while the query planners/optimizers for most modern RDBMS systems are very mature and implement
1577 cost-based optimization through statistics on the data the query planners/optimizers in many Big Data
1578 implementations remain fairly simple and rule-based in nature. While for batch oriented systems this
1579 generally acceptable (since the scale of processing the Big Data in general can be orders of magnitude
1580 more an impact) any attempt to provide interactive response will need very advanced optimizations so
1581 that (at least for queries) only the most likely data to be returned is actually searched. This of course leads
1582 to the single most serious draw back with many of these implementations. Since distributed processing
1583 and storage are essential for achieving scalability these implementations are directly limited by the CAP
1584 (Consistency, Availability, and Partition Tolerance) theorem. Many in fact provide what is generally
1585 referred to as a t-eventual consistency which means that barring any updates to a piece of data all nodes in
1586 the distributed system will eventually return the most recent value. This level of consistency is typically
1587 fine for Data Warehousing applications where data is infrequently updated and updates are generally done
1588 in bulk. However, transaction oriented databases typically require some level of ACID (Atomicity,
1589 Consistency, Isolation, Durability) compliance to insure that all transactions are handled reliably and
1590 conflicts are resolved in a consistent manner. There are a number of both industry and open source

1591 initiatives looking to bring this type of capability to Big Data relational storage frameworks. One
1592 approach is to essentially layer a traditional RDBMS on top of an existing distributed file system
1593 implementation. While vendors claim that this approach means that the overall technology is mature a
1594 great deal of research and implementation experience is needed before the complete performance
1595 characteristics of these implementations are known.

1596 ***Key-Value Storage Models***

1597 Key-value stores are one of the oldest and mature data indexing models. In fact, the principles of key
1598 value stores underpin all the other storage and indexing models. From a Big Data perspective, these stores
1599 effectively represent random access memory models. While the data stored in the values can be arbitrarily
1600 complex in structure all the handling of that complexity must be provided by the application with the
1601 storage implementation often providing back just a pointer to a block of data. Key-Value stores also tend
1602 to work best for 1-1 relationships (e.g., each key relates to a single value) but can also be effective for
1603 keys mapping to lists of homogeneous values. When keys map multiple values of heterogeneous
1604 types/structures or when values from one key need to be joined against values for a different or the same
1605 key then custom application logic is required. It is the requirement for this custom logic that often
1606 prevents Key-Value stores from scaling effectively for certain problems. However, depending on the
1607 problem certain processing architectures can make effective use of distributed key-value stores. Key-
1608 value stores generally deal well with updates when the mapping is one to one and the size/length of the
1609 value data does not change. The ability of key-value stores to handle inserts is generally dependent on the
1610 underlying implementation. Key-value stores also generally require significant effort (either manual or
1611 computational) to deal with changes to the underlying data structure of the values.

1612 Distributed key-value stores are the most frequent implementation utilized in Big Data applications. One
1613 problem that must always be addressed (but is not unique to key-value implementations) is the
1614 distribution of keys across over the space of possible key values. Specifically, keys must be chosen
1615 carefully to avoid skew in the distribution of the data across the cluster. When data is heavily skewed to a
1616 small range it can result in computation hot spots across the cluster if the implementation is attempting to
1617 optimize data locality. If the data is dynamic (new keys being added) for such an implementation then it is
1618 likely that at some point the data will require rebalancing across the cluster. Non-locality optimizing
1619 implementations employ various sorts of hashing, random, or round-robin approaches to data distribution
1620 and don't tend to suffer from skew and hot spots. However, they perform especially poorly on problems
1621 requiring aggregation across the data set.

1622 ***Columnar Storage Models***

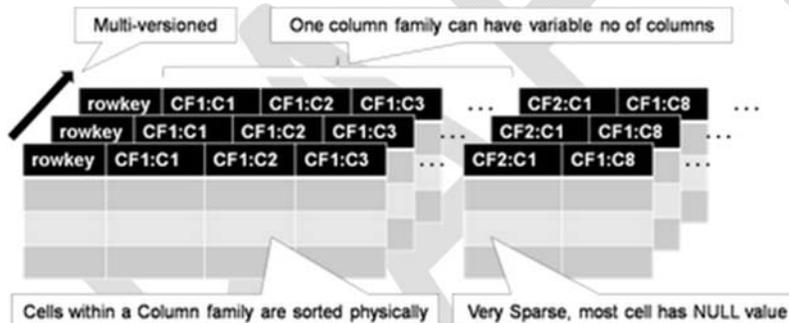
1623 Much of the hype associated with Big Data came with the publication of the Big Table paper in 2006
1624 (Chang, et al., 2006) but column oriented storage models like BigTable are not new to even Big Data and
1625 have been stalwarts of the data warehousing domain for many years. Unlike traditional relational data that
1626 store data by rows of related values, columnar stores organize data in groups of like values. The
1627 difference here is subtle but in relational databases an entire group of columns are tied to some primary-
1628 key (frequently one or more of the columns) to create a record. In columnar, the value of every column is
1629 a key and like column values point to the associated rows. The simplest instance of a columnar store is
1630 little more than a key-value store with the key and value roles reversed. In many ways, columnar data
1631 stores look very similar to indexes in relational databases. Figure 5 below shows the basic differences
1632 between row oriented and column oriented stores.

1633



1634 *Figure B-1: Differences Between Row Oriented and Column Oriented Stores*

1635 In addition, implementations of columnar stores that follow the BigTable model introduce an additional
 1636 level of segmentation beyond the table, row, and column model of the relational model. That is called the
 1637 column family. In those implementations rows have a fixed set of column families but within a column
 1638 family each row can have a variable set of columns. This is illustrated in Figure 6 below.



1639 *Figure B-2: Column Family Segmentation of the Columnar Stores Model*

1640 The key distinction in the implementation of columnar store over relational stores is that data is high de-
 1641 normalized for column stores and that while for relational stores every record contains some value
 1642 (perhaps NULL) for each column, in columnar store the column is only present if there is data for one or
 1643 more rows. This is why many column oriented stores are referred to as sparse storage models. Data for
 1644 each column family is physically stored together on disk sorted by rowkey, column name and timestamp.
 1645 The last (timestamp) is there because the BigTable model also includes the concept of versioning. Every,
 1646 RowKey, Column Family, Column triple is stored with either a system generated or user provided
 1647 Timestamp. This allows users to quickly retrieve the most recent value for a column (the default), the
 1648 specific value for a column by timestamp, or all values for a column. The last is most useful because it
 1649 permits very rapid temporal analysis on data in a column.

1650 Because data for a given column is stored together two key benefits are achieved. First aggregation of the
 1651 data in that column requires only the values for that column to be read. Conversely in a relational system
 1652 the entire row (at least up to the column) needs to be read (which if the row is long and the column at the
 1653 end it could be lots of data). Secondly, updates to a single column do not require the data for the rest of
 1654 the row to be read/written. Also, because all the data in a column is uniform, data can be compressed
 1655 much more efficiently. Often only a single copy of the value for a column is stored followed by the row
 1656 keys where that value exists. And while deletes of an entire column is very efficient, deletes of an entire

1657 record are extremely expensive. This is why historically column oriented stores have been applied to
 1658 OLAP style applications while relational stores were applied to OLTP requirements.

1659 Recently, security has been a major focus of existing column implementations primarily due to the release
 1660 by the National Security Agency (NSA) of its BigTable implementation to the open source community. A
 1661 key advantage of the NSA implementation and other recently announced implementations is the
 1662 availability of security controls at the individual cell level. With these implementations a given user might
 1663 have access to only certain cells in group based potentially based on the value of those or other cells.

1664 There are several very mature distributed column oriented implementations available today from both
 1665 open source groups and commercial foundations. These have been implemented and operational across a
 1666 wide range of businesses and government organizations. Emerging are hybrid capabilities that implement
 1667 relational access methods (e.g., SQL) on top of BigTable/Columnar storage models. In addition, relational
 1668 implementations are adopting columnar oriented physical storage models to provide more efficient access
 1669 for Big Data OLAP like aggregations and analytics.

1670 **Document**

1671 Document storage approaches have been around for some time and popularized by the need to quickly
 1672 search large amounts of unstructured data. Modern, document stores have evolved to include extensive
 1673 search and indexing capabilities for structured data and metadata and why they are often referred to as
 1674 semi-structured data stores. Within a document-oriented data store each “document” encapsulates and
 1675 encodes the metadata, fields, and any other representations of that record. While somewhat analogous to a
 1676 row in a relational table one-reason document stores evolved and have gained in popularity is that most
 1677 implementations do not enforce a fixed or constant schema. While best practices hold that groups of
 1678 documents should be logically related and contain similar data there is no requirement that they be alike
 1679 or that any two documents even contain the same fields. That is one reason that document stores are
 1680 frequently popular for data sets which have sparsely populated fields since there is far less overhead
 1681 normally than traditional RDBMS systems where null value columns in records are actually stored.
 1682 Groups of documents within these types of stores are generally referred to as collections and like key-
 1683 value stores some sort of unique key references each document.

1684 In modern implementations documents can be built of arbitrarily nested structures and can include
 1685 variable length arrays and in some cases executable scripts/code (which has significant security and
 1686 privacy implications). Most document-store implementations also support additional indexes on other
 1687 fields or properties within each document with many implementing specialized index types for sparse
 1688 data, geospatial data, and text.

1689 When modeling data into document-stores the preferred approach is to de-normalize the data as much as
 1690 possible and embed all one-to-one and most one-to-many relationships within a single document. This
 1691 allows for updates to documents to be atomic operations which keep referential integrity between the
 1692 documents. The most common case where references between documents should be use is when there are
 1693 data elements that occur frequently across sets of documents and whose relationship to those documents
 1694 is static. As an example, the publisher of a given book edition does not change and there are far fewer
 1695 publishers than there are books. It would not make sense to embed all the publisher information into each
 1696 book document. Rather the book document would contain a reference to the unique key for the publisher.
 1697 Since for that edition of the book the reference will never change and so there is no danger of loss of
 1698 referential integrity. Thus information about the publisher (address for example) can be updated in a
 1699 single atomic operation the same as the book. Where this information embedded it would need to be
 1700 updated in every book document with that publisher.

1701 In the Big Data realm document stores scale horizontally through the use of partitioning or sharding to
 1702 distribute portions of the collection across multiple nodes. This partitioning can be round-robin based
 1703 insuring an even distribution of data or content/key based so that data locality is maintained for similar

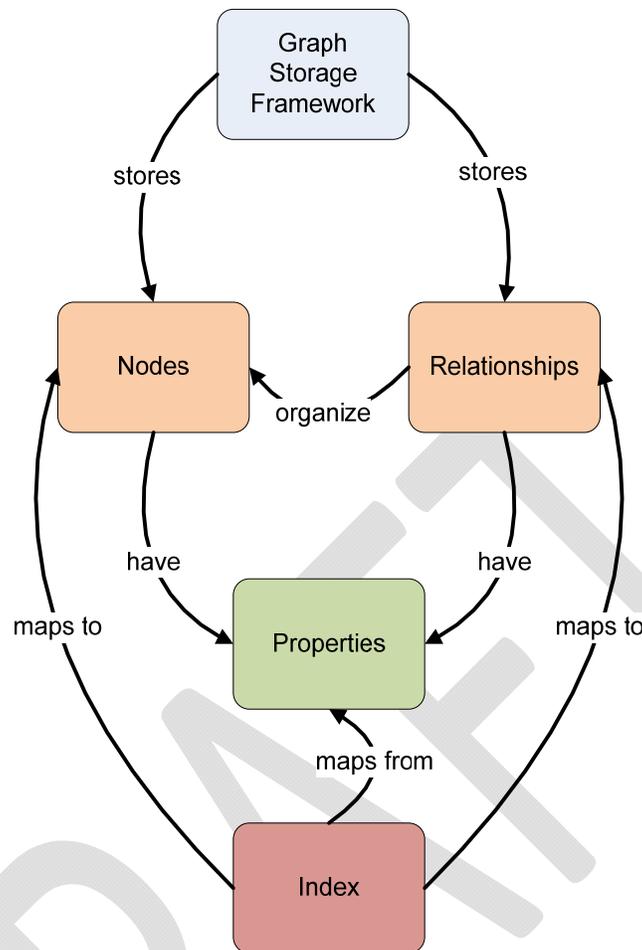
1704 data. Depending on the application required the choice of partitioning key like with any data base can
1705 have significant impacts on performance especially where aggregation functions are concerned.

1706 There are no standard query languages for document store implementations with most using a language
1707 derived from their internal document representation (e.g., JSON, XML).

1708 **Graph**

1709 While social networking sites like Facebook and LinkedIn have certainly driven the visibility of and
1710 evolution of graph stores (and processing as discussed below), graph stores have been a critical part of
1711 many problem domains from military intelligence and counter terrorism to route planning/navigation and
1712 the semantic web for years. Graph stores represent data as a series of nodes, edges, and properties on
1713 those. Analytics against graph stores include very basic shortest path and page ranking to entity
1714 disambiguation and graph matching.

1715 Graph databases typically store two types of objects nodes and relationships as show in Figure 7 below.
1716 Nodes represents objects in the problem domain that are being analyzed be they people, places,
1717 organizations, accounts, or other objects. Relationships describe those objects in the domain relate to each
1718 other. Relationships can be nondirectional/bidirectional but are typically expressed as unidirectional in
1719 order to provide more richness and expressiveness to the relationships. Hence, between two people nodes
1720 where they are father and son there would be two relationships. One “is father of” going from the father
1721 node to the son node, and the other from the son to the father of “is son of”. In addition, nodes and
1722 relationships can have properties or attributes. This is typically descriptive data about the element. For
1723 people it might be name, birthdate, or other descriptive quality. For locations it might be an address or
1724 geospatial coordinate. For a relationship like a phone call it could be the date, time of the call, and the
1725 duration of the call. Within graphs relationships are not always equal or have the same strength. Thus
1726 relationship often has one or more weight, cost, or confidence attributes. A strong relationship between
1727 people might have a high weight because they have known each other for years and communicate every
1728 day. A relationship where two people just met would have a low weight. The distance between nodes (be
1729 it a physical distance or a difficulty) is often expressed as a cost attribute on a relation in order to allow
1730 computation of true shortest paths across a graph. In military intelligence applications, relationships
1731 between nodes in a terrorist or command and control network might only be suspected or have not been
1732 completely verified so those relationships would have confidence attributes. Also, properties on nodes
1733 may also have confidence factors associated with them although in those cases the property can be
1734 decomposed into its own node and tied with a relationship. Graph storage approaches can actually be
1735 viewed as a specialized implementation of a document storage scheme with two types of documents
1736 (nodes and relationships). In addition, one of the most critical elements in analyzing graph data is locating
1737 the node or edge in the graph where the analysis is to begin. To accomplish this most graph databases
1738 implement indexes on the node or edge properties. Unlike, relational and other data storage approaches
1739 most graph databases tend to use artificial/pseudo keys or guides to uniquely identify nodes and edges.
1740 This allows attributes/properties to be easily changed due to both actual changes in the data (someone
1741 changed their name) or as more information is found out (e.g., a better location for some item or event)
1742 without needing to change the pointers two/from relationships.



1743
1744 *Figure B-3: Object Nodes and Relationships of Graph Databases*

1745 The problem with graphs in the Big Data realm is that they grow to be too big to fit into memory on a
 1746 single node and by their typically chaotic nature (few real world graphs follow well defined patterns)
 1747 makes their partitioning for a distributed implementation problematic. While distance between or
 1748 closeness of nodes would seem like a straight forward partitioning approach there are multiple issues
 1749 which must be addressed. First, would be balancing of data. Graphs often tend to have large clusters of
 1750 data very dense in a given area thus leading to essentially imbalances and hot spots in processing. Second,
 1751 no matter how the graph is distributed, there are connections (edges) that will cross the boundaries. That
 1752 typically requires that nodes know about or how to access the data on other nodes and requires inter-node
 1753 data transfer or communication. This makes the choice of processing architectures for graph data
 1754 especially critical. Architectures that do not have inter-node communication/messaging tend not to work
 1755 well for most graph problems. Typically, distributed architectures for processing graphs assign chunks of
 1756 the graph to nodes then the nodes use messaging approaches to communicate changes in the graph or the
 1757 value of certain calculations along a path.

1758 Even small graphs quickly elevate into the realm of Big Data when one is looking for patterns or
 1759 distances across more than one or two degrees of separation between nodes. Depending on the density of
 1760 the graph, this can quickly cause a combinatorial explosion in the number of conditions/patterns that need
 1761 to be tested.

1762 A specialized implementation of a graph store known as the Resource Description Framework (RDF) is
 1763 part of a family of specifications from the World Wide Web Consortium (W3C) that is often directly

1764 associated with Semantic Web and associated concepts. RDF triples as they are known consist of a
1765 Subject (Mr. X), a predicate (lives at), and an object (Mockingbird Lane). Thus a collection of RDF
1766 triples represents and directed labeled graph. The contents of RDF stores are frequently described using
1767 formal ontology languages like OWL or the RDF Schema (RDFS) language, which establish the semantic
1768 meanings and models of the underlying data. To support better horizontal integration (Smith, Malyuta,
1769 Mandirck, Fu, Parent, & Patel, 2012) of heterogeneous data sets extensions to the RDF concept such as
1770 the Data Description Framework (DDF) (Yoakum-Stover & Malyuta, 2008) have been proposed which
1771 add additional types to better support semantic interoperability and analysis.

1772 Graph data stores currently lack any form of standardized APIs or query languages. However, the W3C
1773 has developed the SPARQL query language for RDF which is currently in a recommendation status and
1774 there are several frameworks such as Sesame which are gaining popularity for working with RDF and
1775 other graph oriented data stores.

1776

DRAFT

1777 **Appendix D: Acronyms**

1778	APIs	application programming interface
1779	CIA	confidentiality, integrity, and availability
1780	CRUD	create/read/update/delete
1781	CSP	Cloud Service Provider
1782	DNS	Domain Name Server
1783	GRC	governance, risk, and compliance
1784	HTTP	HyperText Transfer Protocol
1785	I/O processing	Input/Output Processing
1786	IaaS	Infrastructure as a Service
1787	IT	Information Technology
1788	LAN	Local Area Network
1789	MAN	Metropolitan Area Network
1790	NaaS	Network as a Service
1791	NBD-PWG	NIST Big Data Public Working Group
1792	NBDRA	NIST Big Data Reference Architecture
1793	NIST	National Institute of Standards and Technology
1794	PaaS	Platform as a Service
1795	PII	Personally Identifiable Information
1796	SaaS	Software as a Service
1797	SANs	Storage Area Networks
1798	SLAs	Service-level Agreements
1799	VM	Virtual Machine
1800	WAN	Wide Area Network
1801	WiFi	
1802		

1803 **Appendix E: Resources and References**

1804 **GENERAL RESOURCES**

1805 The following resources provide additional information related to Big Data architecture.

1806 [1] Office of the White House Press Secretary, “Obama Administration Unveils “Big Data” Initiative”,
 1807 *White House Press Release* (29 March 2012)
 1808 http://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_press_release_final_2.pdf

1809 [2] White House, “Big Data Across The Federal Government”, 29 March 2012,
 1810 http://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_fact_sheet_final_1.pdf

1811 [3] National Institute of Standards and Technology [NIST], Big Data Workshop, 13 June 2012,
 1812 <http://www.nist.gov/itl/ssd/is/big-data.cfm>

1813 [4] NIST, Big Data Public Working Group, 26 June 2013, <http://bigdatawg.nist.gov>

1814 [5] National Science Foundation, “Big Data R&D Initiative”, June 2012,
 1815 <http://www.nist.gov/itl/ssd/is/upload/NIST-BD-Platforms-05-Big-Data-Wactlar-slides.pdf>

1816 [6] Gartner, “3D Data Management: Controlling Data Volume, Velocity, and Variety”,
 1817 <http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>

1819 [7] Gartner, “The Importance of 'Big Data': A Definition”,
 1820 <http://www.gartner.com/DisplayDocument?id=2057415&ref=clientFriendlyUrl>

1821 [8] Hilbert, Martin and Lopez, Priscilla, “The World’s Technological Capacity to Store, Communicate, and
 1822 Compute Information”, *Science*, 01 April 2011

1823 [9] U.S. Department of Defense, “Reference Architecture Description”, June 2010,
 1824 http://dodcio.defense.gov/Portals/0/Documents/DIEA/Ref_Archi_Description_Final_v1_18Jun10.pdf

1825 [10] Rehtin, Eberhardt, “The Art of Systems Architecting”, CRC Press; 3rd edition, 06 January 2009

1826 [11] ISO/IEC/IEEE 42010 Systems and software engineering — Architecture description, 24 November
 1827 2011, http://www.iso.org/iso/catalogue_detail.htm?csnumber=50508

1828 **DOCUMENT REFERENCES**

¹ The White House Office of Science and Technology Policy, “Big Data is a Big Deal,” *OSTP Blog*, accessed February 21, 2014, <http://www.whitehouse.gov/blog/2012/03/29/big-data-big-deal>.

² “Reference Architecture Description”, U.S. Department of Defense, June 2010.
http://dodcio.defense.gov/Portals/0/Documents/DIEA/Ref_Archi_Description_Final_v1_18Jun10.pdf.

³ Colella, Phillip, Defining software requirements for scientific computing. Slide of 2004 presentation included in David Patterson’s 2005 talk. 2004. <http://www.lanl.gov/orgs/hpc/salishan/salishan2005/davidpatterson.pdf>

⁴ Patterson, David; Yelick, Katherine. Dwarf Mind. A View From Berkeley.
http://view.eecs.berkeley.edu/wiki/Dwarf_Mine

⁵ Leslie G. Valiant, A bridging model for parallel computation, Communications of the ACM, Volume 33 Issue 8, Aug. 1990

⁶ United States Census Bureau, The “72-Year Rule.”
https://www.census.gov/history/www/genealogy/decennial_census_records/the_72_year_rule_1.html. Accessed March 3, 2015.

⁷ NIST SP 500-292, *NIST Cloud Computing Reference Architecture*. September 2011.
http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505

⁸ <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf> (The NIST Definition of Cloud Computing, Special Publication 800-145, September 2011)

DRAFT