



# DESIGN BUILD & MANAGEMENT OF PLANNED BIG DATA ENVIRONMENTS



Alan Senior  
Technical Architect – Data Optimisation  
Computacenter  
[alan.senior@computacenter.com](mailto:alan.senior@computacenter.com)

# Table of Contents

Introduction .....	3
What is Big Data? .....	4
Understanding Big Data .....	4
Large-scale Structured Data Environments (aka Analytics) .....	5
Unstructured, Accumulated data.....	5
Unstructured, Planned data environments.....	6
The Challenge of Planned Big Data Environments.....	7
Core Requirements .....	8
Non-Core Requirements.....	9
An EMC Architecture for Big Data.....	11
Core Storage Hardware – EMC Isilon.....	11
Core Storage Software – Search and Security .....	13
Big Data in the Cloud.....	15
Industry Uses for Planned Big Data .....	16
Conclusion .....	17

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation’s views, processes, or methodologies.

## Introduction

As with any sector in the IT industry, the Data Storage sector likes to label market trends in order to simplify understanding among its customer base. In the past, this was relatively easy and tended toward labelling of specific feature sets:

- Replication
- Snapshots
- Storage Virtualisation
- Data Deduplication

These are some of the more recognisable feature sets that have been sold to customers in the last decade, aimed at meeting customer requirements for faster, larger, more secure transfer and storage of information.

However, in the past few years, the buzzwords used in the IT marketplace have begun to reflect a less easily classified set of challenges. Business Continuity replaced the more explainable Disaster Recovery. Grid Computing, a distribution of compute resources within a localised network, swiftly became Cloud Computing, where the Internet took the place of the local network.

Finally the latest classification—Big Data—which appears at first glance to be a simple concept, but has as yet been unrealised in any but the largest organisations and in the most bespoke of forms.

This article is concerned with Big Data in one of three main forms—the Unstructured, Planned Data Environment. Unlike Structured and Accumulated Big Data, which already have market traction (discussed later in this article) Unstructured Planned Environments have been relatively underplayed in the market, possibly due to size, complexity, and lack of a single product set.

This article looks at the architecture and considerations necessary to create an Unstructured, Planned Big Data solution, particularly with regard to the EMC product set.

## What is Big Data?

Before trying to define Big Data, it is worth considering that:

**Big Data is not a new phenomenon** – While never formally labelled, the tendency for Data Growth vastly beyond the ability of systems and administrators to cope is not new. In the beginning of computing, memory limitations required programmers to use 2-figure dates to avoid the massive costs of additional storage; in the mid/late 90s, Big Data was the tendency for data to increase beyond the bounds of a single server, thus requiring first directly-attached storage arrays, then SAN networks; in the early 2000s, Big Data was any requirement for more than a few tens of terabytes on that SAN. Now, in 2013, Big Data is measured in Petabytes and above. While each time the issue is the same, the solutions to each cycle have been wildly different.

**Big Data is not a single issue** – Big Data is not just a large database, or an accumulation of files. Big Data can be many issues, each with a unique solution, and the storage of the data itself, while a core issue, is far from the end of the matter.

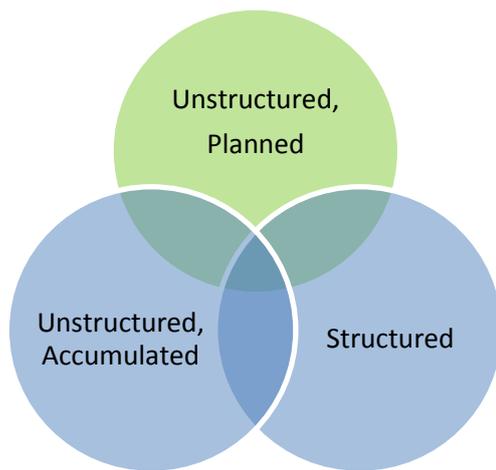
**Big Data is not any one product** – too many vendors will claim that “product x” is the product that will fix your Big Data problems. When tested, it quickly becomes apparent that, at best, a single product will merely form a building block within the larger Big Data solution.

### Understanding Big Data

A large part of the challenge in understanding Big Data appears to be the confusion as to what problem is actually being addressed. The inclusion of several customer challenges within this loose classification has led to this confusion among customers as to the nature of both the Big Data challenge and the solutions required.

At a high level, the Big Data challenges are:

1. Large-scale, structured data environments
2. Unstructured, Accumulated data
3. Unstructured, Planned data environments



**Figure 1: Types of Big Data**

Until now, emphasis has been placed on the first two types of Big Data:

### **Large-scale Structured Data Environments (aka Analytics)**

These environments have grown out of previous solutions such as Data Warehouses where large databases are mined for useful information. As corporate databases have become ever larger and relationships (both internal and external) become more complex, more processing power is required both to extract useful information and to manage the databases instances themselves.

Products for dealing with this type of environment have existed for some time, which tends to exacerbate the confusion around what constitutes a Big Data issue. Many investigators and consultants will immediately assume that a “Big Data” issue is automatically a “Big Database” issue and forget to look at Unstructured Data challenges which may also exist.

### **Unstructured, Accumulated data**

A second type is data which is accumulated automatically through the normal working of IT equipment—log files for servers, network switches, and even bespoke customer equipment (e.g. ATM machines) contain a wealth of information that in past years was thrown away due to the lack of ways to extract useful information.

Recently, powerful extraction tools powered by mechanisms such as Apache Hadoop have allowed organisations to begin making sense of this information. For example, the ability to search through the logs of thousands of ATM machines and compare this with bank account transactions is helping to combat money laundering and other forms of fraud.

This type of use case, while less well developed than the structured data cases, is also developing market traction. Tools such as Splunk are incorporating Hadoop and similar

search algorithms to allow searching of internal and external accumulated data—systems which can scan publically available data and prepare profiles on company employees (both prospective and existing) are already possible.

### **Unstructured, Planned data environments**

Over time, organisations are becoming more reliant than ever on data processing. Entire industries and public services, which a decade ago were still using pen and paper, are now relying more and more on IT for safety, efficiency, and communication.

As organisations accumulate more unstructured file data of all kinds as part of their day-to-day operations, and as this data becomes more valuable it is necessary for organisations to use, store, manage, and retrieve data at a faster rate over vastly increasing volumes.

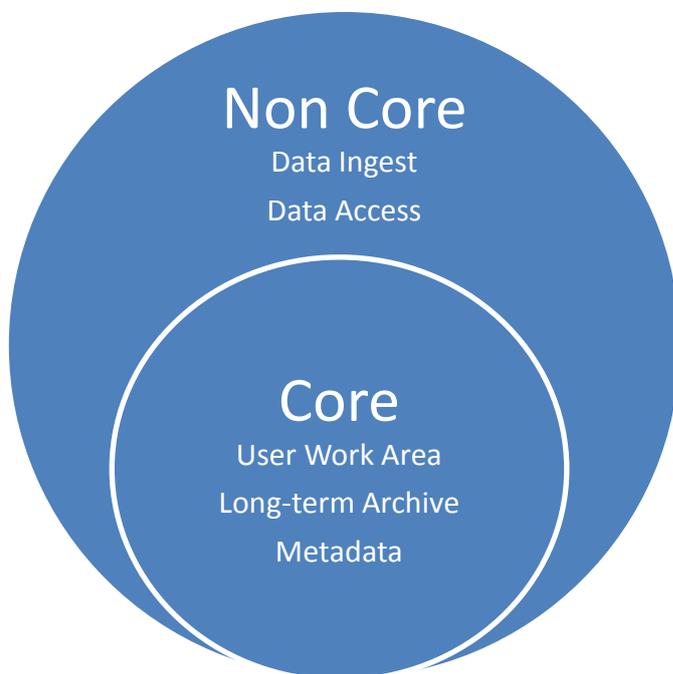
The Planned Data Environment is poorly understood, and tends to be treated as a series of individual challenges; systems are designed to deal with immediate requirements, archiving systems are created separately, and there is little ability to identify and retrieve useful information across the organisation.

Later in this article, we will look at the Planned Data Environment in more detail, identifying the elements which must be considered as part of the overall architecture.

## The Challenge of Planned Big Data Environments

Over time, the requirements of Planned Big Data have become clear due to solutions implemented for existing users of large-scale data. These solutions, while bespoke, all tend to show a set of similar requirements which can be transplanted onto up and coming Big Data users.

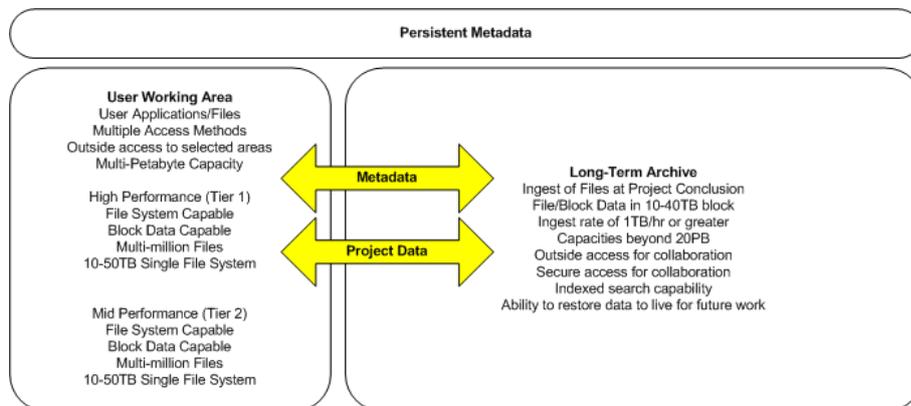
The Big Data systems tend to be split into the core architecture, concentrating on the actual data being held, with non-core requirements tending to be more around movement and protection of the data:



**Figure 2: Core vs. non-Core Requirements**

## Core Requirements

The core requirements are split into three main areas:



**Figure 3: Unstructured, Planned Data Environment**

A **User Working Area**, where live project data can be hosted, allowing applications and direct user access to data. This area is envisaged as being flexible in that it will allow multiple access methods to both file-systems and block data areas. The expectation is that while much of the data held here will require high-performance access (a real-world example given was a 5-10TB file system with 2 million files running against an active application, but up to 50TB per project may be required in some use cases) there may also be a requirement for a lower performance system to hold working projects not requiring tier 1 performance.

While not requiring the capacity of the Long-Term Archive, capacity is still expected to be multiple Petabytes. In a real-world scenario, this area would probably be built first.

A **Long-Term Archive**, where completed project data can be retained, fully indexed, and searchable with data retained for many years. While not expected to run active applications, this area will have to be capable of ingesting the large volumes of data held in the User Working Area.

Outside access to this area could be a key requirement, with emphasis on the security of non-shared data areas.

While existing archiving applications exist, these are not designed to handle the volumes of data that will be held in a Big Data archive (even solutions which promise capacity over 1 Petabyte currently, tend not to have real-world references when challenged).

**Persistent Metadata for Search and Retrieval** - Across both of these areas, it will be necessary to have persistent data relating to each project in the system. In the User Working Area this will be used to identify key facts on a particular project (e.g. data owner, project

duration, requirements for collaboration, retention, etc.). This metadata would then follow the project data into the Long-Term Archive, forming the basis for storage of the data in the long-term (though the metadata itself would need to be stored on a system with enough performance to allow the search to take place).

It is worth noting that while at present metadata can be used to increase the speed of searching, at some point the storage of metadata itself may become uneconomical due to the volumes involved – at this point it may become necessary to search on the data itself (something which is possible at present only on a small scale).

### **Non-Core Requirements**

The Non-Core Requirements are in two parts:

**Ingest** – a large part of the line between a normal file system and the Planned Big Data Environment, is the ability to take in large quantities of data from multiple sources. While a normal file system hold data created specifically by users, the Planned Environment takes data directly from devices. Real-world examples include:

- Gene sequencers
- Video rendering equipment
- Radio telescopes
- Medical imaging equipment
- High volume paper scanners
- GIS mapping systems

These sources and others output a high volume of information which need to be analysed in the User Working Area. However, first it is necessary to bring the information into the core systems. This will require multi-protocol access and may require intermediate systems to provide the middleware between the device and the platform supporting the core systems.

Where metadata for search and security is created at the point of ingest, it must be capable of this creation without significantly slowing the flow of data.

**User Access** – use cases around Big Data are often for both internal and external access (public projects tend to develop the kind of volume requirements that would make a large requirement become a Big Data challenge).

In these cases, it is necessary to develop an access strategy to ensure that levels of access—from internal users, through privileged external users, to the general public—are not able to access restricted information.

This tends to be built into the solution in traditional archive systems but, as previously noted, these traditional archiving systems will either not be usable for the volume of data required, or will require multiple systems, making coherent security policy difficult.

A suggested method for security access is to build a security level into the metadata creation process, ensuring that each piece of data is tagged with a minimum clearance level relevant to the project concerned.

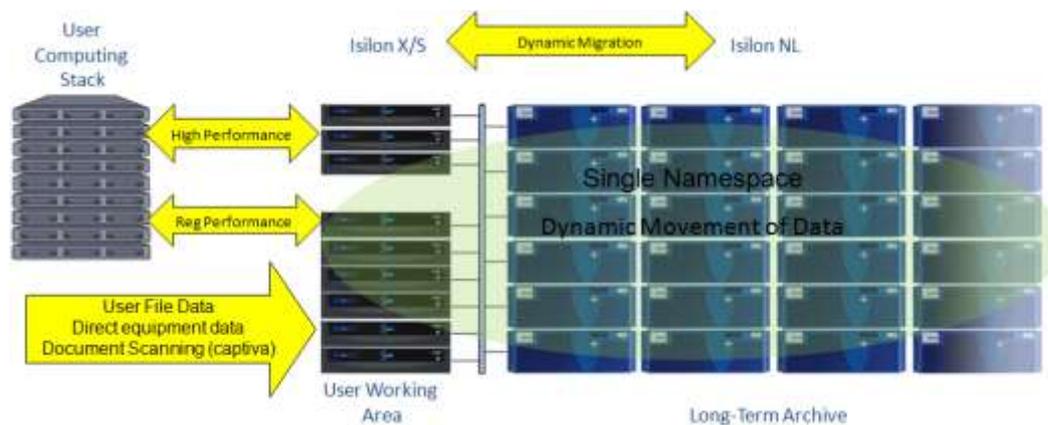
## An EMC Architecture for Big Data

This section maps the points made in previous sections onto the EMC product stack. Where a gap exists between the perceived requirements and EMC strengths, potential solutions are discussed.

### Core Storage Hardware – EMC Isilon

While the general concept architecture sees the User Working Area and Long-term Archive as two separate areas, EMC has a platform which could provide the underlying storage for both. The Isilon® platform has the following strengths<sup>1</sup>:

- Single namespace of up to 20 Petabytes
- Performance scales with capacity – simple building-block architecture
- Data migration from high performance (user working) to low (archive)
- Common connectivity across all tiers
- Multi-protocol (CIFS, NFS, iSCSI) allows for choice of connectivity
- Open protocols allow for custom attachments



**Figure 4: Isilon Ingest to User Working Area, migration to Long-Term Archive**

As shown in Figure 4, the Isilon X and S nodes are used to make up a high and mid-tier User Working Area. In this use case, all ingest is assumed to be live data at least initially, so all data is placed into the User Working Area.

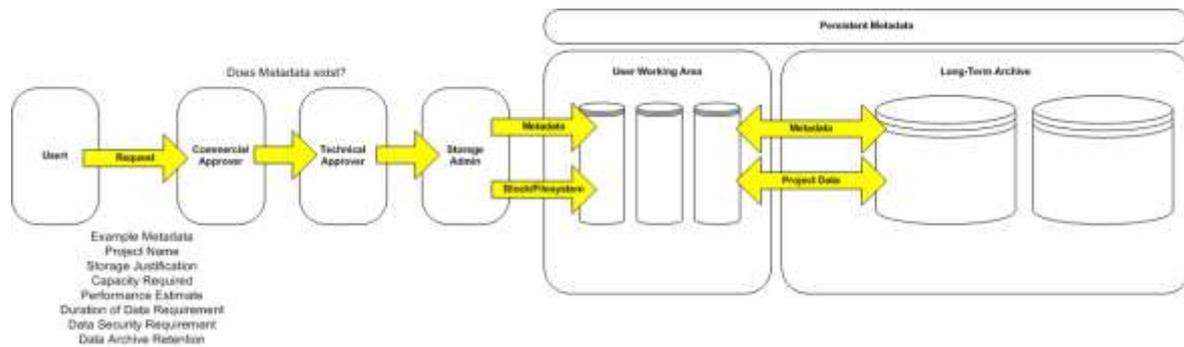
In this case data is shown to come from compute and scanning nodes (i.e. EMC Captiva®). Data could also come direct from devices or via a middleware application or appliance.

---

<sup>1</sup> Source: EMC.com

As data is ingested from multiple sources, the metadata creation process would take place. Initial creation of metadata would be linked to the storage and file-system provisioning process. As part of the request process, as storage is specified for a particular purpose, the key facts for that data area would be gathered as a part of this process.

To illustrate, Figure 5 shows a potential request process for provisioning of storage in the environment:



**Figure 5: Storage Request and Metadata Selection Process Illustration**

At the conclusion of the project, this metadata forms the basis for the creation of the archive volume for the project data – key metrics such as collaboration security, data retention, and information about the project will be picked up by the ingest to archive.

Solution requirement – it should be noted that while Isilon holds metadata as part of the ONEFS® system, this is not user configurable. Metadata creation and management would normally be a 3<sup>rd</sup> party component of the solution.

## Core Storage Software – Search and Security

As discussed previously, the information search function will be crucial to the success of most Planned Data Environments. The purpose of these environments is to provide users with a place to work on large volumes of information, then store and hold that data for many years. When retrieval becomes necessary, it is unlikely the original creator will be performing the search.

In many cases (e.g. research data) it will be vital that all the information on a particular subject be retrieved, not just a single file and multiple original creators may be involved. The search may need to be run multiple times to ensure all relevant information is gathered and this search will be running against vast volume of data, looking for the 0.001% of information that is absolutely required with potentially hundreds of simultaneous searches taking place.

Therefore, after the ability to actually hold and access the volume of data, the search toolkit used in a Planned Data Environment is the most crucial part of the architecture.



Figure 6: Search Stack above User Working Area and Long-Term Archive

EMC currently provides support for Apache Hadoop in various forms but has not yet acquired an intelligent search provider. While Hadoop alone can be used to power a search function, a 3<sup>rd</sup> party tool is generally required to provide human-directed search.

To function correctly in the proposed scenario, the search tool will also need to work with the specific metadata criteria created as a part of the data ingest process. The advantages of this are two-fold; the ability to search specifically against criteria set specifically for a given project, and the ability to direct the search based on security access metadata built into each file.

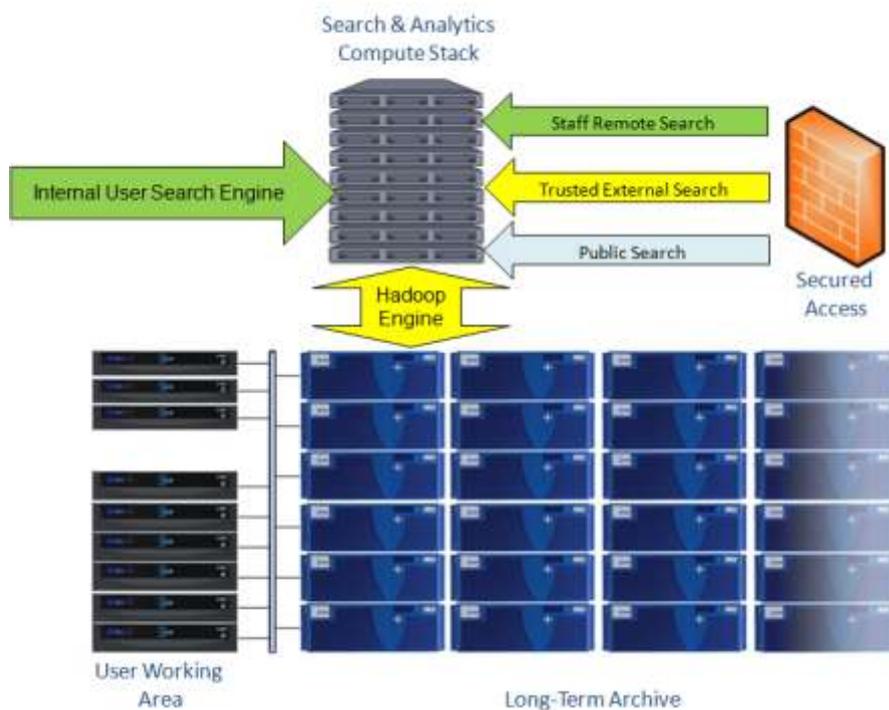


Figure 7: Isilon search stack with security policies

## Big Data in the Cloud

While the current Big Data designs are often driven by the need to collect data within the organisation's data center, it is anticipated that the level of expansion (particularly in the archive) will force customers to consider cloud storage as the logical extension to their Long-Term Archive.

In the proposed scenario, the technology certainly exists to push data from the physical archiving platform; the Cloud Tiering Appliance provides a gateway to a Cloud provider.



**Figure 8: Long-Term Archive - Cloud Expansion**

However, before this can be achieved, the requirements of the Big Data solution must be addressed:

- Metadata needs to be held locally (to be searchable), but must still reference data that has been moved into the cloud
- The ability to recover data from the cloud back to the User Working Area requires a reasonable amount of bandwidth to the cloud provider

A logical upshot of these requirements is that eventually all archived data could be moved to the cloud, retaining only the metadata index to allow for fast searches.

## Industry Uses for Planned Big Data

Possible industry uses envisaged for Unstructured, Planned Big Data Environments are:

**Research Institutions** – large-scale file storage is required to provide a working area where data can be analysed. This could be; Supercomputer outputs, Device outputs (e.g. Gene Sequencers), scanned data for clinical trials, and so on.

Once projects are completed, the final results and raw data need to be stored (raw data storage prevents later claims of tampering with data). If follow-on projects are required, either some or all of the previous data will be required.

**Media & Broadcast** – as greater volumes of both current and past programming are stored in digital format, it becomes necessary to provide systems for long-term storage of edited and unedited footage. High-performance systems provide play-out capability once programmes are recovered from archive.

**Health Providers** – moves to cut costs and increase efficiency in health have led to a reduction in paper records and a move to electronic storage of patient records. The same User Working Area / Long-Term Archive model for Planned Data Environments applies, though the UWA is used to view patient records, and more emphasis is placed on the long-term archive.

**Oil and Gas** – as exploration becomes more difficult, the sophistication of mapping tools and the amount of information provided are rapidly increasing. This information if lost, would take considerable effort and cost to recreate, and requires significant processing power to analyse. Again, the Planned Data Environment model will work, though the search tools may require significant customisation to account for indexing of the unusual data being generated.

## Conclusion

Currently, industry use-cases for Big Data are tending to follow one of three paths:

1. Large-scale structured data environments (Exadata, Greenplum use cases).
2. Unstructured, Accumulated data (Hadoop use cases)
3. Unstructured, Planned data environments (currently, custom-built systems)

This article is concerned with the third type of use-case, Unstructured Planned Big Data Environments, where an organisation's need for data storage is known and understood, but has previously been impossible to meet, mainly due to the requirement for large capacity and high-performance within a reasonable budget.

Up till now, data requirements for more than a few Petabytes in a single system have been successfully met generally by organisations where it is accepted that data is central to the business (e.g. Media, Broadcasters, CERN, NASA, etc.). These organisations have spent large sums over long periods of time to develop bespoke architectures for Big Data.

However, as organisations that are less data-centric are beginning to discover the business value of their data, there is an urgent need to develop architectures for Planned Big Data that are based around mainstream products and standardised architectures, ensuring a reasonable cost and faster time to market than is currently possible.

The article has shown that while the Unstructured Planned Big Data Environment requires more than a single product, the model is achievable using a combination of EMC and 3<sup>rd</sup> party tools.

Based on experience gathered in real world examples, this article looked at the requirements of a use case for an Unstructured, Planned Big Data requirement, including:

- Requirements for Hardware, Software and Deployment
- Requirements for Data in Use (User Working Area) vs Data at Rest (Long-Term Archive)
- Ingest requirement – where does your data come from?
- Data access requirements – who gets access to data?
- Retrieval of Planned Big Data – searching across Petabytes
- Where Big Data meets the Cloud – linking external storage to Planned Big Data systems

The article examined the potential concept architecture built around the industry and use-case, fleshed out further as a product set using the EMC product portfolio where applicable. Finally, the article identified where the concept architecture built around existing use cases can be applied across industries—points of commonality which can be exploited to limit redevelopment time and effort while still ensuring the architecture remains relevant in the new industry.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.