



High Performance Storage System

# HPSS 7.5.1 Upgrade Planning

# Agenda

- HPSS 7.5.1 New Features Overview
- HPSS 7.5.1 Upgrade Overview
- HPSS 7.5.1 Upgrade Planning
- HPSS 7.5.1 New Feature Planning
- HPSS 7.5.1 Upgrade Procedures
- Upgrade Service Options

# HPSS 7.5.1 Overview

- HPSS 7.5.1 New Features
  - Db2 Partitioning
    - Scalability
    - Performance
  - SOID Improvements
    - Improved database performance
  - E2EDI
    - End to End Data Integrity protects the integrity of data at rest
  - TOR
    - Tape Ordered Recall improves tape recall performance
  - Quaid

# HPSS 7.5.1 New Feature Overview

- Db2 Partitions

- Advantages

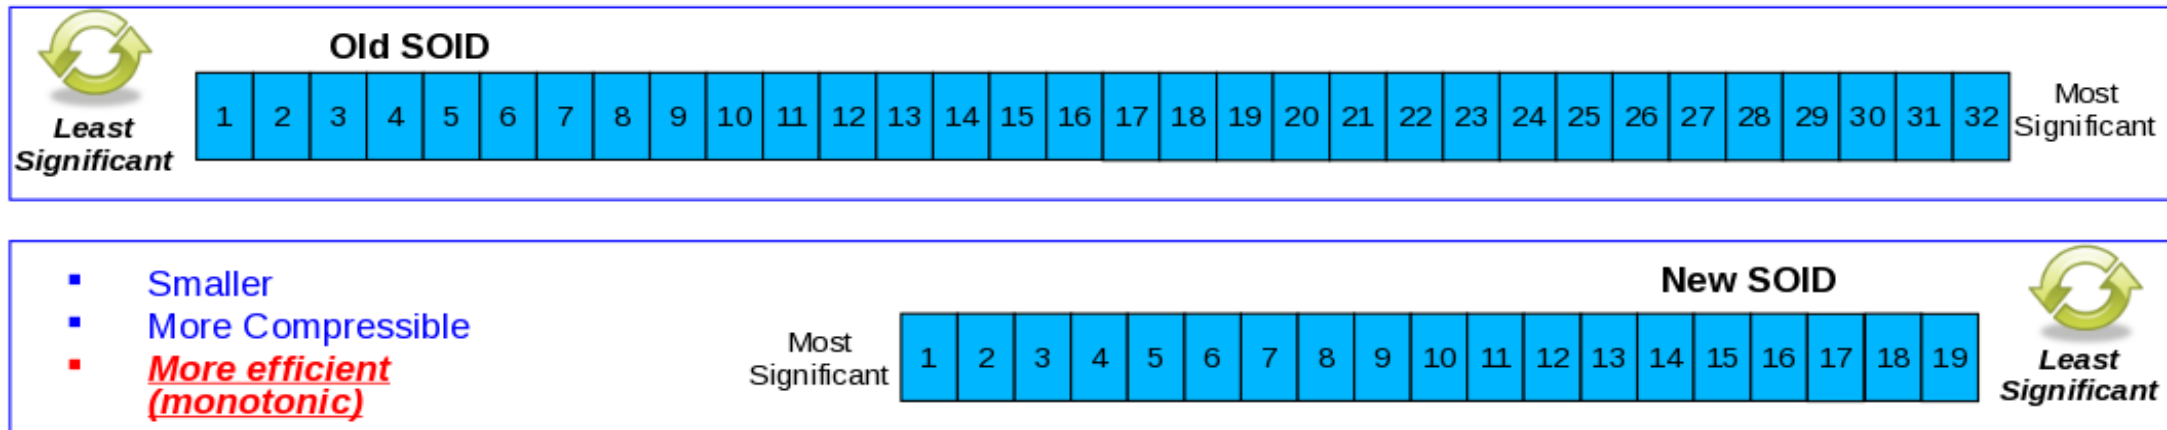
- Increased metadata transactional performance
      - Transactions are not bottlenecked by a single logging CPU/thread
      - Database performance can scale to 100% of hardware performance
    - Increased Performance for Analytic Operations (Db2 History)
    - Increased performance of administrative functions
      - Db2 maintenance and backups of each partition can be done in parallel so administrative functions on large databases can complete faster
    - Foundational component for exascale computing
      - Foundation for off-host Db2 and distributed database over multiple servers

- Considerations

- Additional configuration components required
    - Use Log Shipping for High Availability

# HPSS 7.5.1 New Feature Overview

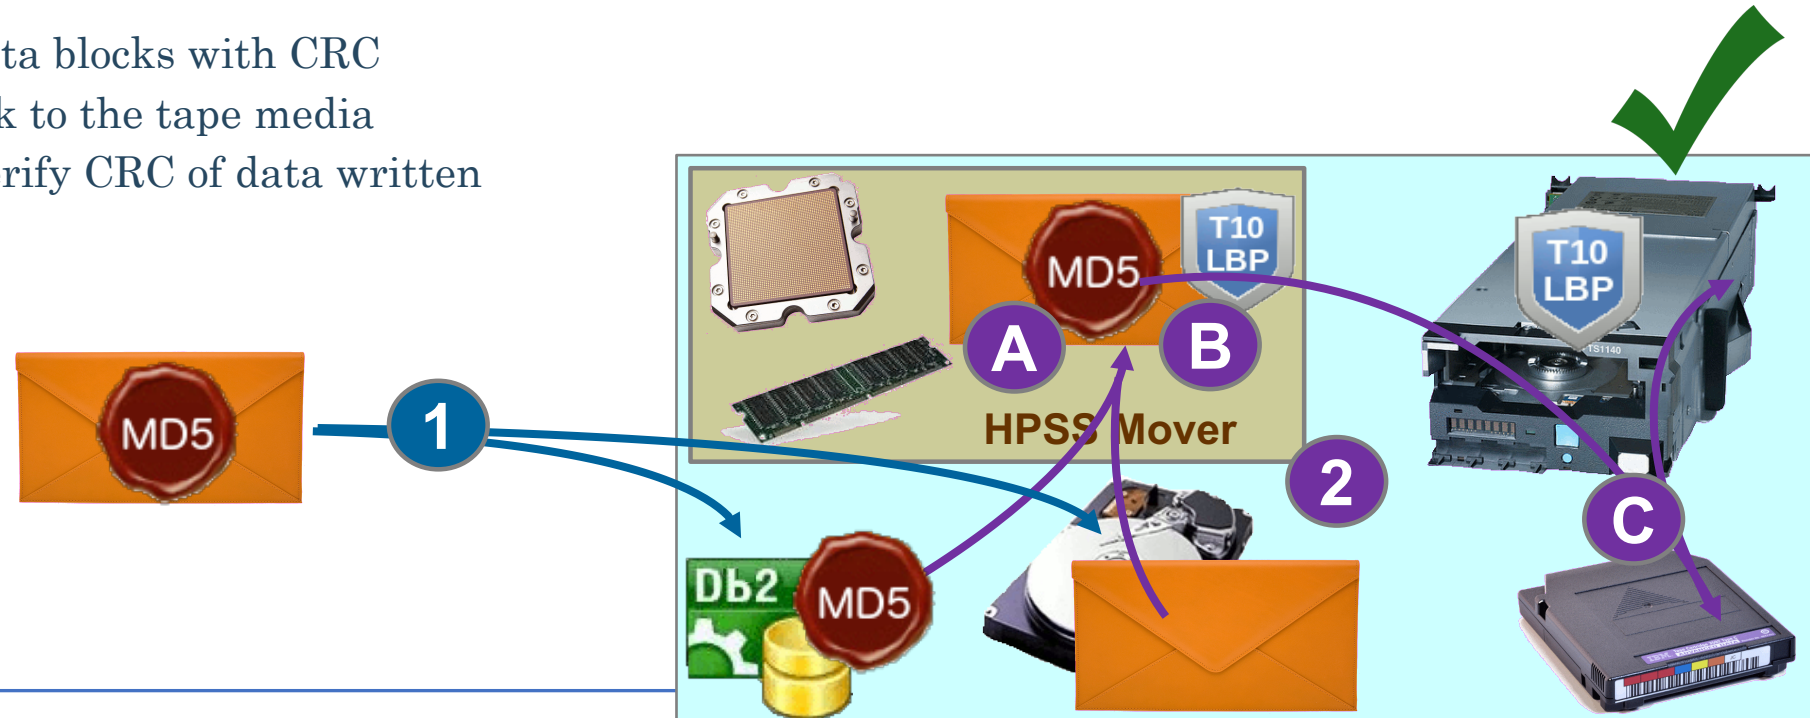
- SOID Improvements
  - Smaller SOID value (reduced from 32 to 19 bytes)
  - SOID structure reorganized to improve index performance and compression



# HPSS 7.5.1 New Feature Overview

- HPSS End-To-End Data Integrity

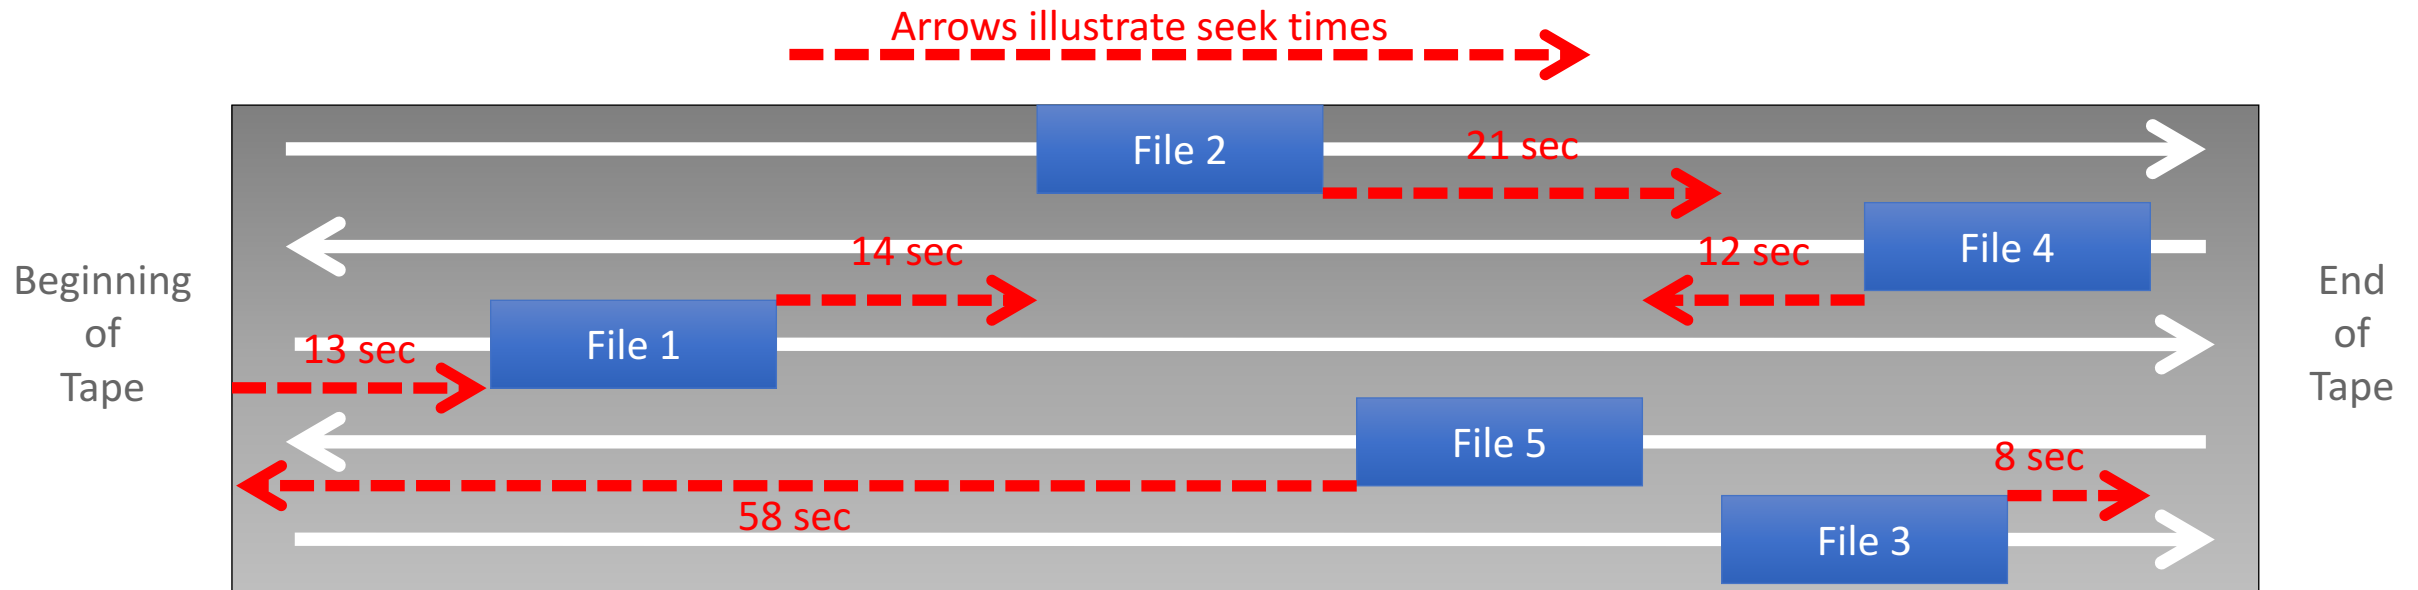
- 1) User sends file to HPSS disk and file checksum is stored in Db2
- 2) On disk-to-tape migration, HPSS Tape Mover will:
  - A. Read the file data from HPSS disk and calculate file checksum
  - B. Interlace the CRC for each tape block being written to tape
  - C. The tape drive
    - ✓ Receives the data blocks with CRC
    - ✓ Writes the block to the tape media
    - ✓ Immediately verify CRC of data written



# HPSS 7.5.1 New Feature Overview

- HPSS Tape Ordered Recall

- Random tape recalls are slow when they are not properly ordered
- Ordering tape recalls by tape offset (LTO) may result in shorter recall times
- Enterprise tape and HPSS support 'Recommended Access Order' (RAO)
- RAO tape recalls on Enterprise tape can take 30% to 60% less time



# HPSS 7.5.1 New Feature Overview

- Quaid
  - New tool that can sort files by volume and stage files
  - Uses new Batch Stage API to stage files in the background
  - Pass Quaid a file list and it will stage files to disk cache
  - File stage status is sent to stdout for processing
  - Supports Multiple Input Methods
    - Command Line
    - Input File
    - Xquery
  - Supports Multiple Filters
    - COS, File Family, UDA, Account, Volume
    - Filter Operations (equal, not equal, etc)
  - Example scripts are available from Support



# HPSS 7.5.1 Feature Overview

- Improved Segment Unlink Performance
- Batch Resource Create/Delete Improvements
- Retry Offline Disks
- I/O Aware Mounts
- API pkg-config Support
- LBP Verify Tool

# Upgrade Overview

- Upgrade Support
  - IBM will work with the sites to assess the best upgrade options for a successful upgrade
  - IBM will help determine best Q-Rep Options
  - IBM will work with sites to determine the best Partition configuration
  - IBM will provide a detailed conversion guide for sites to use during the conversion
- Detailed Conversion Guide for HPSS 7.5.1
  - Instructions for Pre-Req installation
  - Database configuration instructions
  - Q-Rep setup and execution instructions
  - Q-Verify setup and execution instructions

# Upgrade Overview

- Conversion and Q-Rep Tested and Ready
  - Many hours of development and testing during LLNL upgrade
  - LLNL Upgraded 3 Systems – 2 Production and 1 Pre-Production
  - LANL Upgraded 4 Systems – 2 Production and 2 Test
  - Over a billion rows in some tables at both sites
  - Undisclosed Customer Upgrade completed
- Q-Verify
  - Validates all Target data and changes during the conversion

# Upgrade Overview

- Three Conversion Paths
  - HPSS 7.4.2p1 -> HPSS 7.5.1
  - HPSS 7.4.3p2 -> HPSS 7.5.1
  - HPSS 7.4.3p3 -> HPSS 7.5.1
- Success Starts with a Good Upgrade Plan
  - Plan for all changes to the system
    - Hardware changes
    - Software changes
  - Plan for Testing
    - Test Applications
    - Test Conversions

# HPSS 7.5.1 Upgrade Planning

- Upgrade Decisions
  - Hardware
    - Migrate to New Hardware
    - In Place Migration (Single Server)
  - Pre-Requisites
    - RHEL 6.8/6.9
    - RHEL 7.4
  - Q-Replication Options
    - In Place (Single Server)
    - Server to Server
    - Offline
  - Partitioning

# HPSS 7.5.1 Upgrade Planning

- Red Hat Operating System Options

- RHEL 6.8

- Upgrade Servers to RHEL 6.8/6.9
    - RPM upgrade
    - Minimal Disruption to the server environment
    - Best option for in place upgrade

- RHEL 7.4

- Install RHEL 7.4 on servers
    - Overwrites RHEL 6.x configuration and data for in place upgrade
    - Major Disruption to server environment, like installing new servers
    - Good option when migrating to new hardware



# HPSS 7.5.1 Upgrade Planning

- Upgrade Logistics
  - Q-Rep only conversions
  - Are there local site conversion steps required?
    - Storing checksum information in the comments field?
- Q-Rep Pre-Requisites
  - Db2 10.5 FP8
  - Db2 Federation (may require Db2 restart if not already configured)
  - Db2 Q-Rep Software
  - MQ software
  - HPSS Q-Rep python scripts and templates

# HPSS 7.5.1 Upgrade Planning

- Q-Rep Options
  - Online – Replication is done while HPSS is running
    - Q-Rep metadata related downtime is minimal
  - Offline – Replication is done while HPSS is offline or down
  - In Place
    - Q-Rep Source and Target databases are on the same server
    - Metadata space requires 2x the current database size, current usage should be under 40%
    - Online conversion will have 2x increase in processing, memory, and logging load on the server
  - Server to Server
    - Q-Rep Target database is on a different server from the Source database
    - High Speed Network connection between Source and Target servers
    - Off loads much of the processing to the Target server during Online conversion



# HPSS 7.5.1 Upgrade Planning

- Are multiple Db2 Partitions needed?
  - How many files in 5 years?
    - Greater than 500 Million files
  - Other HPSS actions that drive database transactions?
    - Large Number of Recalls
    - Analytics Workloads
  - How large is the database?
    - Partitioning can help parallelize Db2 maintenance and backups for large databases
  - How is the metadata storage configured?
    - SSD, Flash
    - Disk
    - Metadata storage performance must be able to support partitioned database environment
  - Is HA or failover needed?
    - Use Db2 Log Shipping

# Upgrade Planning


- Upgrade Information
  - Review and Document current HPSS System Configuration with Support
    - lshpss
    - lsnode output for servers
  - Review current Db2 configuration with Support
    - lsdb2 output
    - Automatic Storage Settings
    - Db2 Federation
    - Database Performance Information
      - Memory Usage
      - Db2 Log LUN utilization
      - CPU Utilization on Core Server
      - Database Storage Used

# Upgrade Planning

- Recommended Upgrade Paths
  - New Hardware/Core Server Upgrades
    - Q-Rep Server to Server and RHEL 7.4 are best options
  - In Place Upgrades
    - Q-Rep In Place and RHEL minor upgrade are best options
- Test Planning
  - Upgrade a test system first and start testing early
  - Test application compatibility with HPSS 7.5.1
  - Build and Test new HPSS features into applications
  - Test system performance if possible
  - Test Q-Rep process and dry-run conversion steps

# HPSS 7.5.1 New Feature Planning

- Tape Ordered Recall (TOR)
  - Enterprise drives
    - Supports Recommended Access Order (RAO)
  - LTO Drives
    - Logical Tape Order only
  - New Batch Stage API
  - New Staging Best Practices
    - Use background stage commands instead of foreground stages
    - Send many stage requests in parallel rather than one at a time
    - May require tuning Core Server parameters for Copy Operations
    - Work with Support on Tuning Core Server
  - Evaluate using Quaid to stage files



# HPSS 7.5.1 New Feature Planning

- End to End Data Integrity
  - Logical Block Protection (LBP)
    - CRC is generated for each tape block written to tape
  - File Hash
    - Algorithm Selection could impact migration performance
      - Use algorithms that can use cryptographic features on CPU
    - Will the client generate a file hash to pass to HPSS?
      - Save the File Hash by setting the File Attribute
  - Tape Mover CPU requirements
    - Migration requires 2 CPU cores for File Hash and LBP calculation
    - Stage requires 1 CPU core for LBP calculation
- pkg-config Scripts
  - Makefile.macros no longer needed by client API programs
  - Use pkg-config scripts when compiling Client API code
  - See the HPSS Programmer's Reference for tips on how to use pkg-config



# Upgrade Procedure Overview

- Pre-Upgrade Planning
  - Review System Configuration with Support
  - Document Upgrade Plan
- Upgrade Test System
  - Verify and Test Upgrade Procedures using test system metadata
  - Test HPSS 7.5.1 with site applications and new features
- Dry Run Production Metadata Conversion
  - Convert production metadata on test system or new hardware
  - Validate conversion procedure and timing
  - Verify production metadata conversion results
- Upgrade Production System

# Upgrade Procedure Detail

- Create New 7.5.1 Databases
- Q-Rep Setup
- Q-Rep Build and Start
- Q-Verify Build and Start
- Monitor Conversion
- Complete Q-Rep
- Downtime Upgrade Tasks
- Post Conversion Tasks

# Upgrade Support and Services

- Manage HPSS 7.5.1 upgrades through the PCR process
- IBM Support PM will work with the client to assess:
  - Scope of the upgrade
  - Responsibilities for services and deliverables
- PCR will be generated based on the outcome of the assessment
- PCR may not result in additional costs if IBM Support Staff workload is not increased beyond standard support consultation
- IBM Managed Upgrade provides an option to procure additional support hours for IBM to manage and execute the upgrade procedures and testing



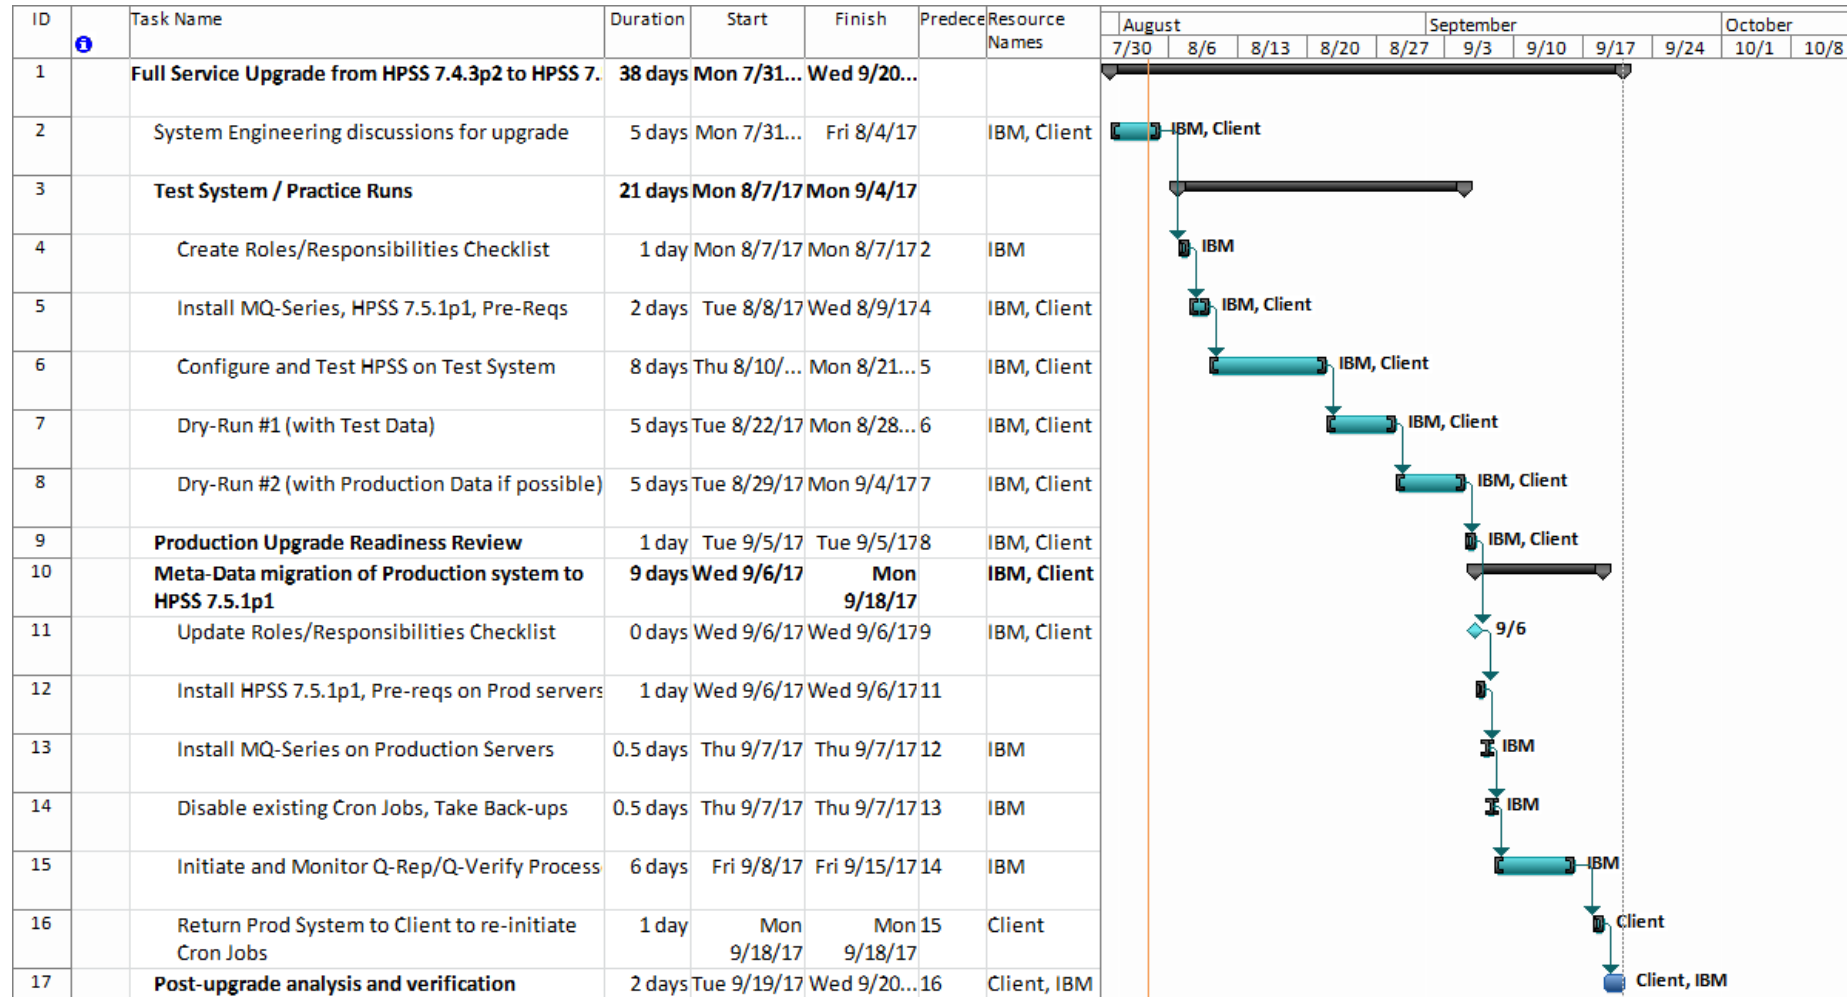
# Roles and Responsibilities

- IBM Managed Upgrade
  - IBM
    - Review Db2 and HPSS Configuration
    - Review Upgrade Plan
    - Provide Required Software and Scripts
    - Install Software and Scripts
    - Provide Conversion Procedure
    - Execute Conversion Procedures on Test System
    - Review Test Results
    - Execute Conversion Procedures on Production System
  - Client
    - Provide Configurations for Review
    - Document Upgrade and Test Plan
    - Execute Test Plan on Test System

# IBM Managed Upgrade



High Performance Storage System



# Conclusion

- 3 Sites currently working on Upgrade to 7.5.1
- 7 Sites have started planning an Upgrade to 7.5.1
- When are you planning your upgrade to 7.5.1?