Information Technology Supporting Documentation
Commonwealth of Pennsylvania
Governor's Office of Administration/Office for Information Technology

Document Number:	**OPD-INF001B**
Document Title:	**Database Management Systems: Production and Operational Standards**
Issued by:	**Deputy Secretary for Information Technology**
Date Issued: February 23, 2005	**Date Revised: November 18, 2010**
Domain:	**Information**
Discipline:	**Data Management**
Technology Area:	**Data Management Systems**
Referenced by:	**ITP-INF001**
Revision History Date:	**Description:**
04/16/2009	**Added paragraph 2.2**
11/18/2010	**ITP Refresh**

Production and Operational Standards	**Rationale**
1.0 Design Specifications and Requirements	
1.1 A high-availability strategy such as failover, mirroring, and/or the use of online backups are to be used for databases requiring 24/7/365 availability.	Provides for uninterrupted data access.
1.2 Online Transaction Processing (OLTP) data that is used by online users for mission-critical day-to-day operations and Online Analytical Processing (OLAP) data used for ad-hoc queries and decision support is to be segregated into separate databases and/or servers.	Reduces the performance impact on production applications caused by large queries and ad-hoc processing demands common in decision support systems.
1.3 Large databases that have high availability and currency requirements, support high transaction rates, and share data across multiple agencies are good candidates for centralization.	These types of databases require a great deal of operational and database administration support.
1.4 All replicated data is to be designed to be read-only. Updates are to occur through the source where the data originates to facilitate ease of data management. Replication of data is to be based on needs such as availability, security, performance, or decision support.	Data quality and integrity are more manageable when replicated and distributed data is read-only. Replication can: - ensure uninterrupted access to critical data. - isolate production data from external users. - facilitate load balancing through synchronization of distributed databases.
1.5 New databases are to use a relational Database Management System (DBMS) and support ANSI Standard SQL (currently SQL92), and are to also comply with Open Database Connectivity (ODBC) and Java Database Connectivity (JDBC) standards.	ANSI S3.135-1992 Database Language SQL as delimited by FIPS PUB 127-3. - Nonstandard language features are to be used only when the needed operation or function cannot reasonably be implemented with the standard features alone. The use of standard features is to be favored over optional vendor-specific features to prevent being locked into a vendor-specific solution.

***** Predefined supplemental document type codes are listed below:
APP = Appendix **BPD** = Best Practice Document **GEN** = General Information Document
OPD = Operations Document **RFD** = Existing Supporting Document Referenced by this ITP **WHP** = White paper

Production and Operational Standards	Rationale
1.6 The source data that populates a data warehouse is to be accurate. Updates are to occur only through the OLTP source where the data originates. End users are to have read-only access.	Data needs to be accurate to ensure good business decisions.
1.7 Systems are to be designed to reject invalid data and to assist end users with data entry corrections. All updates to an authoritative source OLTP database are to occur using the business rules that own the data, not by direct access to the database. Commercial Off-The-Shelf (COTS) software is to be configurable to allow for data validation.	If data quality validations are built into new and existing application systems, it reduces the risk of inaccurate or misleading data and reduces the need for data cleansing.
1.8 Effective data modeling practices are to be an integral part of all application development projects.	Specific roles and responsibilities are to be established for Data and Database Administrators to ensure that business rule requirements, and both the logical and physical database design requirements, are met.
1.9 Purge criteria is to be established for all databases and is to be in accordance with agency record retention schedules.	Data storage is expensive. Data that is no longer needed is to be purged or archived to a less expensive media in accordance with data retention requirements, policies, or historical significance.
1.10 Metadata is to be captured and maintained for both OLTP and OLAP environments to facilitate data sharing within an agency and among agencies.	It is easier and more cost effective to capture metadata when new applications are being developed. Business process owners and developers, including contractors, are responsible for documenting elements stored in the database. Standards for metadata will be documented as part of Data Standards.
2.0 Production Specifications and Requirements	
2.1 When implementing a data warehouse or data mart, a network assessment is to be done to determine the potential impact on the network.	Ad-hoc workloads can be unpredictable and can place significant demands on the network.
2.2 Desktop database including, but not limited to, Oracle, IBM DB2, Microsoft SQL Server, and MySQL. They are not to be used for enterprise, mission-critical, or multi-user solutions.	These types of systems are difficult to support and are not to be used to support critical business functions.
2.3 The proliferation of Microsoft Access Databases and other desktop databases to develop stand-alone applications is discouraged. Microsoft Access is suitable as a front-end reporting tool when the reports users are limited to a small group or division, but Access is not to be used as an enterprise reporting tool or as a shared interface to enterprise database solutions.	These types of systems are difficult to support and are not to be used to support critical business functions.
2.4 Microsoft Access Databases are not to be used as enterprise or multi-user solutions.	Microsoft Access is marketed as and only appropriate as a desk-top solution. Performance degrades significantly when multiple users attempt simultaneous access. Being a file-based database, as opposed to transactional-based, auditing capabilities are severely limited,

Production and Operational Standards	Rationale
	transactional and point-in-time recovery is not supported, and there is no administrative support. File (database) and data security features are also lacking.
2.5 Production databases supporting mission-critical applications are to be recoverable to a point-in-time and point-of-failure. Logging of database activity is to always be active for mission-critical production applications. Database copies are to be sent offsite weekly, at a minimum.	A plan is to be in place and tested at least twice each year to prevent loss of data due to application/software errors, or from hardware failures. Database backups are to be scheduled frequently enough to ensure optimum recovery times.
2.6 An internal auditing process is to be put in place to ensure that the agency is in compliance with all security, privacy, and information dissemination laws.	Database Administrators (DBAs) are to work closely with security officers to ensure that information is stored, managed, accessed, and secured in compliance with Health Insurance Portability and Accountability Act (HIPAA), Internal Revenue Service (IRS), and Criminal History Record Information Act (CHRIA) Privacy laws and other laws impacting agency program areas.
2.7 All platforms used for hosting mission-critical applications are to be fully supported by the vendor and provide optimum performance. DBMS vendors acknowledge these platforms as preferred or tier-1.	This provides better access to the latest versions of products, timely patches, and to knowledgeable technical support on preferred/tier-1 platforms.
2.8 Database backups for mission-critical applications are to be tested at least twice each year. An updated Disaster Recovery (DR) Plan is to be available that documents the operational procedures required for the recovery, including procedures for retrieving off-site copies of database backups.	Backups cannot be assumed to be complete until they are tested. Errors in scripts, backup software, or other software products may prevent recovery. This is an opportunity to ensure that all DBAs understand the steps needed to perform a recovery, and testing will ensure minimal downtime when an actual recovery needs to be performed.
2.9 Databases for mission-critical applications are to be monitored proactively for capacity planning purposes and are to maintain high availability. Statistics such as transaction rates, allocated extent size, system CPU, and archive log volume, are to be included.	A database is never to stop functioning for foreseeable events. Events such as file systems filling up, objects not able to allocate additional extents, and objects reaching maximum extents, are to be actively monitored and a DBA notified before the problem causes the database to halt.
2.10 Database permissions are to be granted at the minimum level required. Limit the members of the system or database administrators' role to trusted database administrators. Create custom database roles, if required, for better control over permissions. Application programs or interfaces are never to be given *sysadmin/sysdba* authority. Default accounts are to be changed. Production passwords are to be changed from test or development environments.	Employees or contractors who are granted the *sysadmin/sysdba* role have the authority to do anything in the DBMS. Likewise, application programs or interfaces could be modified to perform unauthorized updates.

Production and Operational Standards	Rationale
A process is to be put in place to monitor the activity of those with System or Database Administration authority.	
2.11 Error logs and event logs for security-related alerts/errors are to be constantly monitored. The use of Intrusion Detection Systems (IDS), especially on mission-critical online database servers, is to be considered and implemented if they are vulnerable.	IDS can constantly analyze the inbound network traffic, look for trends, and detect Denial-of-Service (DoS) attacks and port scans. IDS can be configured to alert the administrators upon detecting a particular trend.
2.12 DBAs are to remain current with the information on the latest service packs and security patches released by DBMS vendors. All the service packs and patches are to be carefully evaluated and applied according to vendor recommendations.	Known vulnerabilities can easily be exploited by hackers.
2.13 Application Capacity Planning is to be done at least once a year to determine the impact on mission-critical databases. Physical resource utilization monitoring is to be employed for trend analysis in support of the capacity plan.	Capacity planning will enable system administrators and management to accurately forecast future hardware/software capacity requirements. The plan is to address topics such as transaction rates, physical storage needs, software licensing, and machine utilization, and is to be updated periodically to anticipate future needs at three-, six-, and twelve-month intervals or longer, based on procurement lead time.
2.14 A process is to be established to regularly stress test mission-critical applications.	Stress and/or load tests are required prior to each major application release and as work loads increase. The stress test will simulate anticipated normal and peak usage of the application and the impact on the infrastructure. The results of the test are to be reviewed by a performance team comprised of the application, database, network, and other technical support personnel responsible for performance issues.
2.15 Statistics related to normal database operations, networks, and applications are to be collected and available for a minimum of three months.	Having a baseline set of statistics encompassing server, database, network, and application utilization will prove to be invaluable, when problem issues arise.
2.16 A Change Management Process to coordinate and communicate infrastructure, application, and database-related changes is to be put in place.	Having a defined process for approval and notification provides overall coordination of activities and limits exposure to multiple simultaneous changes from occurring. This becomes very valuable when problem situations arise; changes that were made to the infrastructure, applications, or databases can be identified easily.
2.17 DBMSs and the operating systems from which they operate are to be on version/release levels that are fully supported by the vendor.	Unsupported software is no longer being updated to fix newly discovered security vulnerabilities or other problems that occur due to environmental changes.