Shared IT Infrastructure

August 2012

OFFICE OF INNOVATION & TECHNOLOGY

CITY OF PHILADELPHIA

Table of Contents

Executive Summary

The purpose of this document is to guide agencies toward leveraging existing shared IT infrastructure, services, processes and support staff in order to minimize risk and lower the overall cost of IT efforts.

This document focuses on the existing shared infrastructure used by multiple agencies and is not a complete listing of every product used by every agency.

The City of Philadelphia's (the City's) Shared IT Infrastructure has been built to support this vision. It is a standardized environment that currently supports enterprise computer systems within and across agency boundaries. The infrastructure is designed to accommodate growth and replacement of hardware, middleware, software and communications as new business needs arise or when efficiencies can be realized by upgrading or replacing existing components.

It is also the intent of Executive Order No. 12-11 to ensure that the City has the most effective information technology infrastructure, one that meets the needs of City departments, boards, commissions and agencies and those that they serve. This includes but is not limited to:

- Identifying the most effective approach for implementing new information technology directions throughout City government;
- Improving the value of the City's technology assets and the return on the City's technology investments;
- Ensuring data and information system security and continuity;
- Planning for continuing operations in the event of disruption of information or communication services;
- Supporting accountable, efficient and effective government by every City department, board, commission and agency.

This document is intended to provide sufficient technical detail regarding the various components of the City's Shared IT Infrastructure and, in Appendix 1, denotes the level of support and investment the City has made in specific products and technologies. While continually evolving, it is based on industry standard open system solutions that provide a high degree of vendor neutrality, maximum flexibility, and the agility needed to meet the ever-growing service delivery needs of the City's Executive Branch. The use of open standards is critical to the City's ability to interact with constituents and business partners across the Internet. The focus on specific products and technologies is equally important in order to minimize the staffing resources needed to support a shared, consolidated infrastructure.

The organization of this document is based on the IT Architecture Stack depicted below, where each layer represents a set of technologies put in place to support specific business processes. At every layer, the products and technologies implemented were selected to maximize investment dollars and to ensure architectural integrity (i.e., Product A works with Product B). This architecture stack is currently used to deliver information and services to every major user community in State government.

Executive Summary (cont.)

Specific benefits of the architecture include:

- *Reduced costs for new applications*
- *Improved access to legacy data*
- *Centralized help desk, backup and recovery services*
- *Faster delivery of applications across a multitude of devices and networks*
- *Minimized data redundancy through data sharing*
- *Reduced dependency on proprietary components*
- *Reduced risk in reliable operations, security and change management*
- *Expert staff specially trained on enterprise platforms*

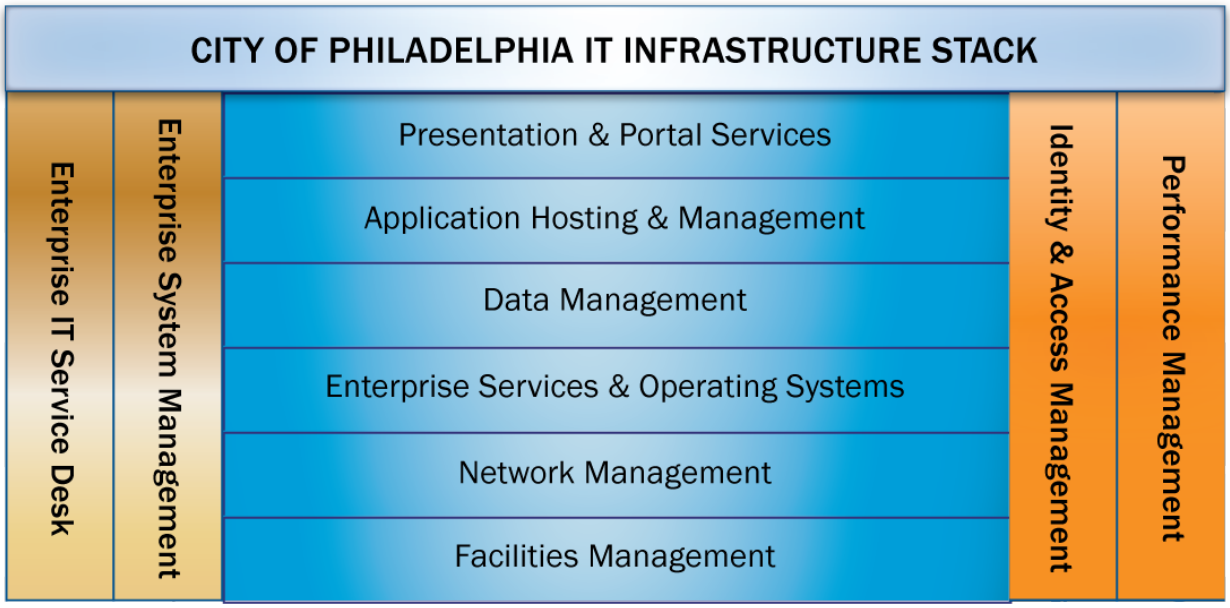

While the existing infrastructure is described by way of an architectural stack, the City is in the processes of undertaking an Enterprise Architecture program to focus on the Business, Information, and Technology needs of the City as an enterprise. This program will help achieve success in the government-to-business, government-to-constituent, and government-to-government domains. This initiative is intended to bring together executives and IT leaders from all City agencies to define the common vision for accomplishing this mission.

Facilities Management

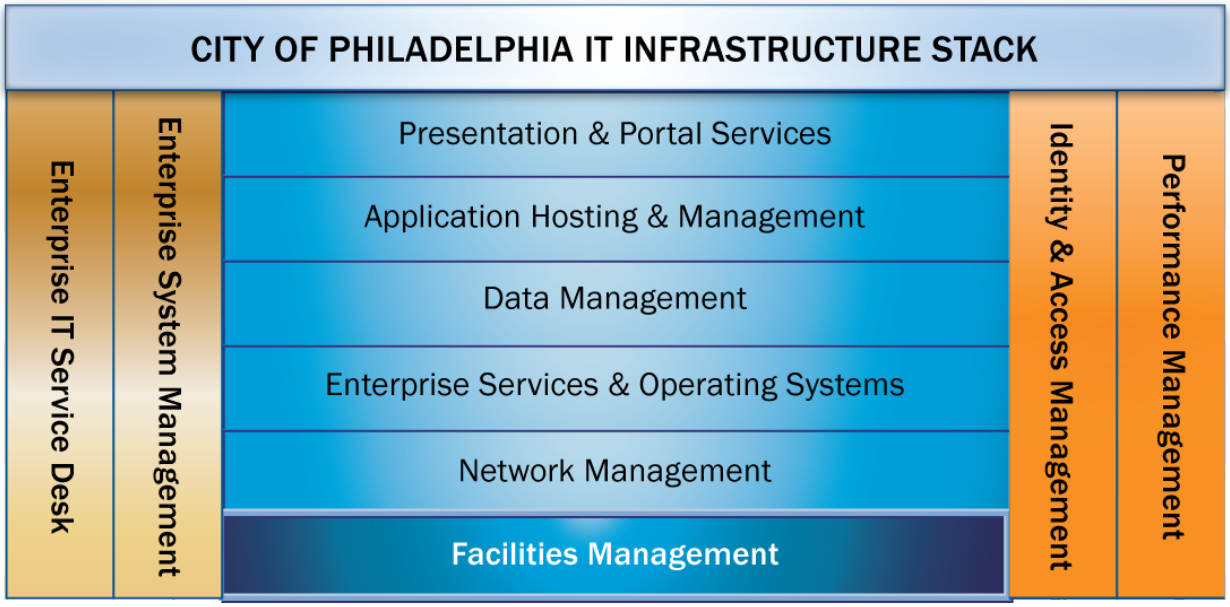

The data center facility provides backup and recovery services for the mainframe environments and critical infrastructure services and serves as City's disaster recovery facility. The data center facility will be leveraged as a production-hosting environment to support applications in conjunction with the two data centers.

Plans are also underway to provide agencies with alternative geographic locations where mission critical applications can be hosted in the event of a disaster scenario at the primary facilities or as a means to provide additional capacity.

Physical Security

In addition to the secure office building location of the data center facilities, OIT also employs additional layers of physical security to ensure that client assets are safe, secure, and protected against outside intrusion and unauthorized access.

Building Security

Uniformed and civilian personnel control the movement of all persons within the office buildings. Access to secured areas is permitted via an authorized badge access system that is maintained by the OIT security team. Security Cameras are placed strategically throughout the data center facilities to prevent against unauthorized access or tampering activity.

Cabinet Systems

The majority of the servers are housed within standard cabinet systems. Access is limited to authorized system administrators to perform standard software, hardware, and diagnostic services.

Facilities Management (cont.)

Control Center

Operation of the primary data center is managed by a central Operations Center. This control center is manned by a group of support professionals twenty-four hours a day, three hundred and sixty-five days a year. The responsibility of Operations Center personnel is to ensure the availability, reliability, and operational status of all production servers, the network, the environmental systems, and security systems within the facility. Facility Management and Network Monitoring systems and software are utilized by Control Center personnel to proactively monitor and display the status of these systems within the facility.

Alarms

Alarms are strategically placed throughout each data center facility and within the server rooms to alert personnel in the event of an environmental system failure or fire.

Electrical

The goal is for each data center to have redundant power systems in order to achieve maximum availability and reliability of all systems. Operations Center personnel closely monitor internal power distribution systems to maximize system uptime.

Commercial Power

The two data centers facilities are fed by two separate power grids, providing greater resiliency in electrical availability.

Power Distribution

A network of Power Distribution Units (PDUs) and Panels that distribute and supply power to all critical servers and associated equipment is housed in each respective facility. Servers equipped with redundant power supplies are cross-connected to PDUs and panels that are connected a single UPS bus. This arrangement provides sufficient power redundancy to enable critical servers and other equipment with dual power supplies to remain up and operational in the event of a PDU or panel failure.

Uninterruptible Power Sources

Each data center utilizes and maintains multiple Uninterruptible Power Sources (UPS) that allow all critical systems and associated equipment to remain powered up and operational in the event of a power failure. All critical equipment at each facility is connected to a two phase UPS Backup System which engages automatically when primary and secondary commercial power feeds fail. These systems include both battery and diesel generated backup power.

Environmental Climate Control

Each data center is equipped with a complete environmental system to guarantee optimal cooling and humidity levels in order to facilitate the availability, reliability, and continued operation of all systems. Operations Center personnel monitor these environmental system controls. Each facility has N + 1 Redundant Liebert units ducted together to provide the environmental climate control to keep all systems and associated equipment operational and within the prescribed temperature and humidity limit boundaries. Any abnormal environmental climate conditions are immediately logged and reported to the OIT Facilities Group for resolution as well as vendor notification where necessary.

Facilities Management (cont.)

Fire Detection and Suppression Systems

Each data center has a complete fire detection and suppression system equipped with an annunciator panel that shows the current status of the fire detection and suppression system. The Operations Center personnel proactively monitor these panels. The fire suppression system dispenses a fire retardant gas that extinguishes fire immediately upon detection.

Network Management

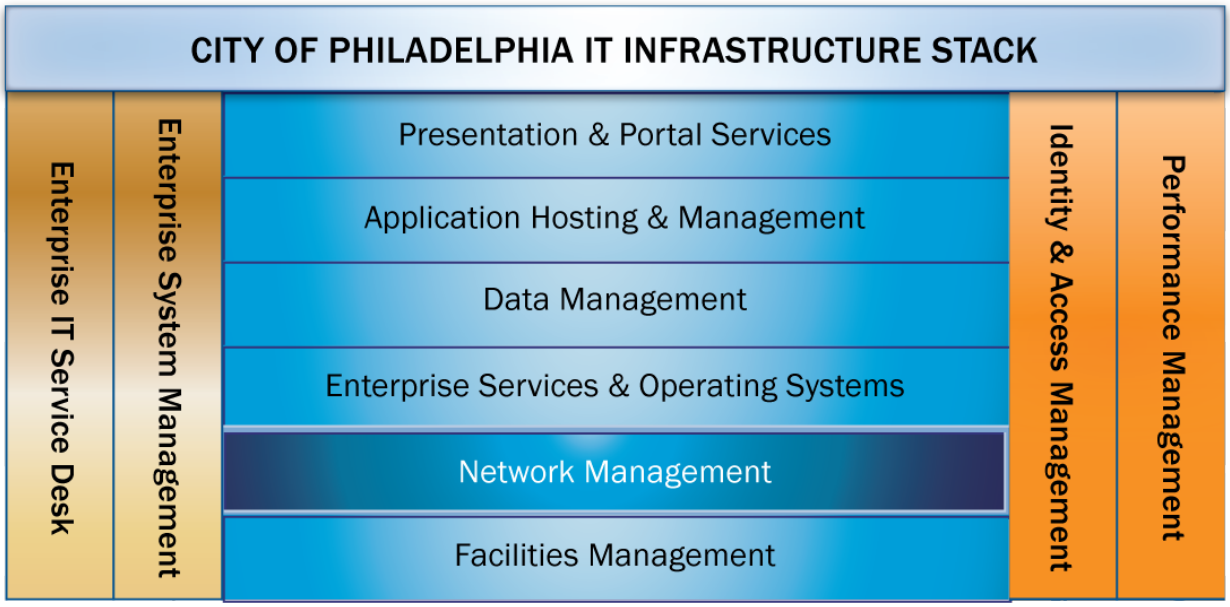

Network Architecture

The City of Philadelphia Office of Innovation & Technology implements, manages and maintains heterogeneous network infrastructure, providing WAN access and aggregation, remote access, backbone, data center, including access to E-government and IP based mainframe application services and Internet access services. This is in support of the operational requirements of City of Philadelphia (City) Executive Branch departments, boards, commissions and agencies, City Council, First Judicial District of Pennsylvania, and the Philadelphia Office of the District Attorney as well as providing secured access to publically accessible City hosted resident, business, and informational services applications.

The City's network currently supports over 20,000 IP addressable devices. Included in this device count are over 1,000 routers/switches and security appliances, approximately 1,000 data circuits and over 200+ application servers.

Legacy Carrier Grade Network Architecture

The City's network (CityNet) provides carrier grade backbone and remote facility (local access) services to the City of Philadelphia (City) Executive Branch departments, boards, commissions and agencies, City Council, First Judicial District of Pennsylvania, and the Philadelphia Office of the District Attorney. The City's network is a diverse, multi-protocol environment providing both dedicated and switched services in support of centrally hosted (City data centers) enterprise E-government and mainframe based application services and distributed intranet applications and internal business services.

The CityNet is comprised of thirteen main node facilities. The node facilities provide aggregation services for remote traffic and facilitate carrier-to-carrier or network-to-network interfaces utilizing the City's Optera duel ring core infrastructure. The currently contracted carrier services supporting the legacy CityNet are provided by Verizon. The backbone is designed with multiple, redundant paths to

Network Management (cont.)

increase service reliability and availability while maintaining the isolation of City department, board, commission and agency traffic across the backbone. Primary transport technologies serving the legacy backbone are T1, T-3, OC3, OC12, SONET, TLS and DWDM. Remote facilities connect to their central nodes or to the CityNet node facilities primarily with T-1, frame relay, or point-to-point services. The Inter-LATA traffic aggregation is supported via Verizon OC3, OC12 or T3 technologies. Intra-LATA transport services are provided by Verizon using OC12 and DS3 technologies.

The Next Generation Services Network Architecture

The impetus for the development of the Next Generation Services Network was to capitalize on the potential synergies available through governmental consolidation by leveraging available infrastructure assets and to develop a standard enterprise model for providing essential networking services City-wide, support for industry standard technologies such as 1 and 10 Gigabit Ethernet, support for end-to-end Quality of Services to support IP based VOIP/Telephony, and Video Conferencing initiatives.

Through the utilization of City of Philadelphia owned dark fiber assets, the vision of building a Citywide fiber based network with protected on-ring presence in each of the City of Philadelphia's data centers:, and core network locations has been realized and fulfilled. With the two of the major ring components completed (East and West) and the proposed expansion of the fiber infrastructure (North and South). NGSN is now positioned to add significant supplied carrier and converged IP services to support Executive Branch operations, public safety initiatives and critical strategic objectives set forth by the Mayor.

NGSN Strategic Benefits

The Next Generation Services Network provides expanded, on ring points of presence to deliver network access services to the City of Philadelphia. The NGSN is comprised of 11 main node facilities located in each of the communication LATAs within the City. The NGSN provides a significant increase in bandwidth capacity and support for IP based services in comparison to the legacy City network. The NGSN core, interconnecting the City's data center facilities: The data center currently provides 2.5 gigabits capacity on protected fiber rings. The NGSN will support full convergence of networking services, data, voice and video, end-to-end Quality of Service (QOS) and Private Virtual Cloud Services.

NGSN Ongoing Migration Strategy utilizing Carrier Ethernet Services

OIT has co-located the NGSN optical ring with the Carrier Ethernet Networks in each communication LATA and are using OIT managed Multiprotocol Label Switching (MPLS) technologies to seamlessly provision the Carrier Ethernet Services for City departments, boards, commissions and agencies. The utilization of Carrier Ethernet to support remote client facilities enables the City's MPLS Services to provide path isolation through the use of L2 and L3 virtualization to support isolation of department, board, commission and agency traffic and to rapidly provision bandwidth to support increasing capacity and IP services demands. Multi-tenant facilities supported through Carrier Ethernet Services enable OIT to implement a shared services model utilizing a protected Ethernet circuit, OIT managed router and switch, multiple sub-interface configurations to support traffic isolation and individual tenant capacity demands over the MPLS enabled backbone.

Network Management (cont.)

Legacy Network Migration Strategy

City departments, boards, commissions and agencies not positioned to migrate to MPLS supported Carrier Ethernet Services require OIT to deploy Carrier OC at NGSN Nodes to support aggregation of legacy ILEC Frame Relay (FRASI, Frame to ATM) and ATM network to network (NNI) circuits. OIT provider edge routers will provision the Cross-LATA ATM links over the City's MPLS backbone.

Advanced Services Supported on the Next Generation Services Network

Real-time Voice and Video Applications Voice Gateway/Dial Tone/Call Manager Video and teleconferencing network isolation Secure Guest access Wireless Network Isolation Robust Data Center interconnects

Internet Services

Current Internet Services (circuits) are contractually provided to the City by Cavalier. Migrated to two 1 GB aggregated Ethernet circuits in 2011

Internet Gateway and Perimeter Controls

OIT has designed and implemented the architecture to provide services to monitor, access, report, track, and manage external access. This environment addresses security for inbound and outbound risks at the lowest total cost of ownership. This environment provides:

- Web Content Filtering
- Intrusion Prevention
- Quality of Service (QoS)
- Access Control

N-Tier Internet Architecture

The City of Philadelphia supports a multi-tiered environment in which to host E-government applications. The n-tier environment provides secure, but direct access to the City of Philadelphia informational and critical line of business application systems. Current security policy dictates that web access directly from the Public Internet is limited to externally facing web servers or content filtering systems. The n-tier environment supports presentation, business logic and data layers. The data center-hosting environment has recently undergone a complete refresh, replacing all core layer 2 and layer 3 components, including new MDF and IDF distribution facilities.

Enhancements to the E-government/data center-hosting environment include:

- Redundancy at all Network layers
- Redundant network connections for all servers at the OIT data center
- Redundant power grids
- Increased throughput
- Access Policy Enforcement
- Integrated firewall service modules in Fail-Over Configuration
- Multiple Security Zone support 2-tier, and 3-tier
- Intrusion Detection and Prevention Systems, Monitoring and Logging
- Network Services Distribution Model
- Simplified Cable Management for servers, SAN, KVM, IP DRAC

Network Management (cont.)

Tunneling, simple pass-through proxy, 'double tier hops', and other techniques that do not apply policy or process to an inbound communication at each tier, are not allowed -to do so would compromise the integrity of all remaining applications that follow the security policy.

Secure Remote User Access

The City maintains several mechanisms to provide secure remote user access to resources:

- For applications that do not meet the traditional E-government model for web, presentation and data layer design, extranet connectivity is available.
- Extranet connections require point-to-point connections from the extranet partner to the extranet infrastructure either via a point-to-point data circuit, or through an IPSec tunnel across the Internet. The cost of these connections varies based on the type of data circuit ordered, and the equipment required to terminate the circuits. The Remote Access VPN solution provides SSL and IPSEC VPN services to City and non-City users.

Network Systems Management

Real-time proactive monitoring is performed on all OIT managed network infrastructure devices, both Wide Area Networks (CityNet) and the data center core and E-government hosting environments. All devices are monitored for performance, availability and health statistics, including CPU and memory usage, operational temperature control, fans, power supply, individual interface and sub-interface events, Routing Protocol status and other hardware and software event.

Enterprise Services & Operating Systems

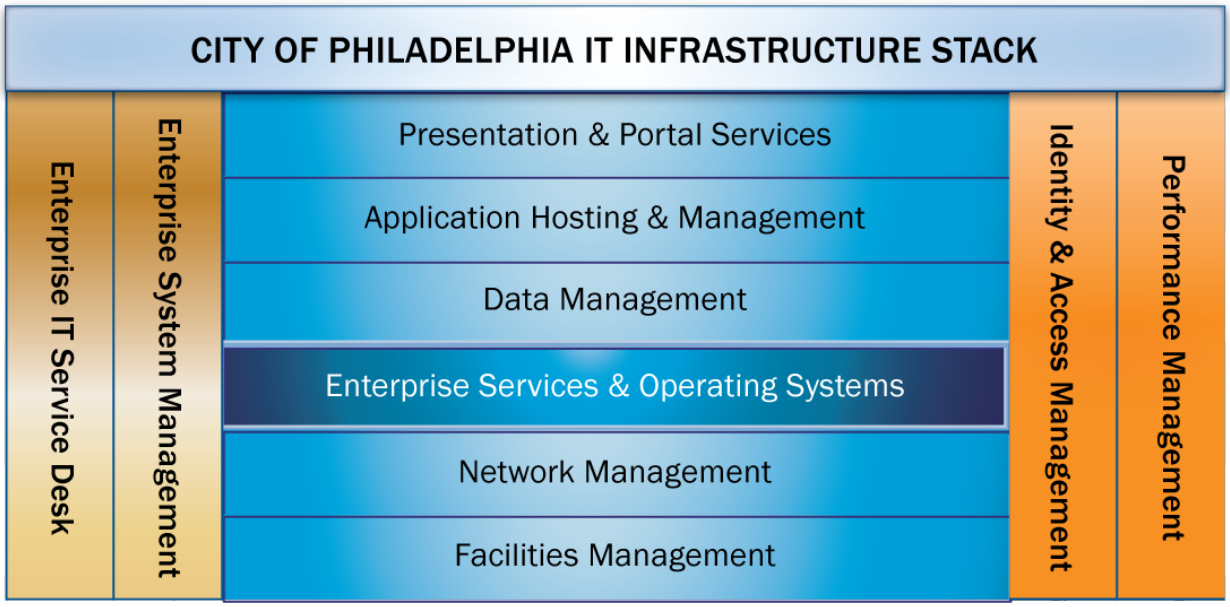

Shared Server Infrastructure

Mainframes and servers are centralized to offer a common location to manage the distributed environment. Cabinets are provided to rack servers and eliminate excess footprint. Implementation of a standard KVM (Keyboard, Video, Mouse) matrix switching backbone solution at both facilities has improved floor space utilization, cable management and server access as well as reduced equipment requirements and power consumption. Optimizing key server resources through common logical and physical environments positions the City to properly plan, manage and control a growing server infrastructure. For all servers housed in this environment, OIT and the agency may share the administration of the solution components.

Based on the best-supported environments by the IT community, the SSI supports the following operating system platforms:

- IBM z/OS
- IBM AIX
- Sun Solaris
- Linux
- Microsoft Windows

Server Virtualization and Consolidation

Another key data center optimization strategy pursued by OIT is server virtualization and consolidation. Implementing this strategy is dependent on technological advances in both hardware and software that have now cut across all operating system platforms noted above. This approach saves on data center floor space, power, and cooling per unit of processing capacity. In addition, operations, administration, and maintenance can be addressed more efficiently and less expensively. Consequently, for new applications,

Enterprise Services & Operating Systems (cont.)

OIT is driving the deployment of virtualized servers as the preferred approach. For existing applications, OIT is pursuing server virtualization and consolidation where it makes business sense to do so (e.g., at the point of equipment refresh or maintenance renewal).

OIT is also pursuing virtualization and consolidation of infrastructure services as more and more agencies leverage the City's enterprise hosting architecture. Two specific examples are given below.

Enterprise eMail Services

The Office of Innovation & Technology maintains a highly available, redundant enterprise infrastructure to facilitate inbound and outbound email processing for City agencies. Gateway services include message routing, anti-virus and anti-spam scanning.

All inbound and outbound emails are scanned at the gateways for virus content. Anti-spam processing is also available, on an opt-in basis for City agencies.

The City is in the process of consolidating to one messaging platform – Microsoft Exchange. This consolidation will create a centralized Active Directory Resource Forest to support a citywide messaging and calendar platform based on Exchange Server 2010 including the necessary systems to monitor and manage the new environment.

Storage Area Network

The City manages a Storage Area Network (SAN). Storage Management offers fully redundant storage arrays, with over 1.2 PB of storage currently in use. The SAN consists of a redundant core to edge fibre channel communication that provides physical connections, a management layer that organizes the connections and storage layer that controls data delivery and security. Storage devices are connected to servers in a networked fashion, using directors to build the topology. The City uses a variety of storage array types to optimize performance and minimize price based on storage needs.

The SAN currently supports connection speeds of 1,2 and 4 GB. Upgrades are in the process to take this to 8 GB in the next year.

In order for a server to connect to the SAN, a Host Bus Adapter (HBA) must be installed in the server. Two HBAs can be used to provide redundant paths to the SAN; this mitigates the risk of having a single point of failure. Once connected, disk space can be allocated from the storage array(s) and dedicated to a server. SAN technology presents many benefits to server data storage, such as:

- Centralized storage management
- Ability to add disk capacity dynamically
- Ability to replace a deficient server without loss of data
- Faster response time than internal SCSI disks
- Potential for improved backup and disaster recovery techniques
- Better storage attributes – hardware RAID, dynamic sparing, remote data copy, mirroring, and more.

Storage Management also offers boot from SAN. Using this method, all OS drives are replicated to the OARS recovery site for quicker server recovery.

Enterprise Services & Operating Systems (cont.)

Backup and Restore Services

OIT Storage Management is currently converting to CommVault for backup and restore services available to clients within the multiple security zones.

Clients consist of Windows, Linux, Solaris, AIX, Novell and VMware systems, as well as Oracle, SQL, and DB2 databases. Other clients are available upon request.

These services require CommVault software loaded on the target server that selects the data for backup on the server, and then sends the selected data to the CommVault server by the way of TCP/IP.

Storage Management requires the creation of a User ID with Root/Administrator authority on the target server, which is used to install the client, monitor backups, and troubleshoot any problems that may occur during daily backup processing.

For all servers at the City's data center facilities, an additional Network Interface Card (NIC) should be installed and connected to the Storage Management Backup Network in order to reduce the backup window, and eliminate network contention.

Note that these services are for backup and restore of the server data only. Data archiving is a different process.

Basic Server Backup/Restore Policy (Unstructured Data)

The standard client backup begins at midnight 0000 hours (12:00 AM) with backup duration dependent upon client hardware and network bandwidth. Most clients are usually finished the backup processing by 0600 hours (6:00 AM), and must be completed by 0730 hours (7:30 AM).

The first backup is that of a full system, meaning that every file not specifically excluded by the CommVault configurations files is sent to the backup server. Subsequent full backups are done every 12 weeks. Incremental backups will be done in between so that only those files that have changed since the last backup are sent to the backup server. This method reduces network bandwidth consumption and backup storage requirements. Every 4 weeks a Synthetic Full will be created. A Synthetic Full creates a new tape merging the full with all the incremental taken since the last full. Backup data is stored on virtual tape.

The standard backup policy will retain unstructured data for a period of 60 days.

Structured Data Services

CommVault will fail to backup files that are open for writing. For data that must be available to an application 24/7, CommVault provides other clients that must be utilized.

Oracle Database Backups

Storage Management uses the Oracle Recovery Manager (RMAN) in conjunction with CommVault to backup and restore Oracle database instances. In most cases, clients depend on 24/7 operations that cannot be interrupted for backup processing. Storage Management utilizes a hot backup procedure that allows database operations to continue while the database is backed up unless the client has specified otherwise.

Enterprise Services & Operating Systems (cont.)

The standard RMAN hot backup policy consists of a full backup of the entire database once per week. Backups are also performed on a nightly and non-cumulative incremental basis for the remainder of the week.

Control files are backed up nightly along with the full and incremental database backups. Archive Logs are also backed up via the nightly RMAN scripts unless the logs are managed by the migration client. Oracle parameter files, password files, and other configuration files are not managed by RMAN, and should be backed up using the CommVault server.

All RMAN backups are tracked by an RMAN recovery catalog residing on the backup server, and all backup pieces generated by RMAN are stored by CommVault. Storage Management offers a recovery window of twenty-one (21) days for the standard Oracle client.

This means that Storage Management keeps all RMAN backups necessary to restore a database to a point in time equal to twenty-one (21) days prior to the current time or today minus twenty-one (21) days. Once an RMAN backup piece is no longer useful for this recovery window, it is expired and no longer available for restore operations.

Exchange Services

A client add-on for CommVault can be installed and configured on each server running Microsoft Exchange Server.

A full "hot" backup is done on Sunday evening at 1930 hours (7:30 PM) on each CommVault client. The "hot" backup is an open-file-supported-backup of each active Exchange instance.

Incremental "hot" backups are performed from Saturday through Thursday at 1930 hours (7:30 PM). Incremental backups include all files that have been changed since the last full backup.

SharePoint Services

This section is currently under development.

Rebuild and Restore

The most direct rebuild and restore method is to rebuild the server, at which point Storage Management can reload the CommVault Client and restore the data. This will require the most time to return to the operational state.

Data Management

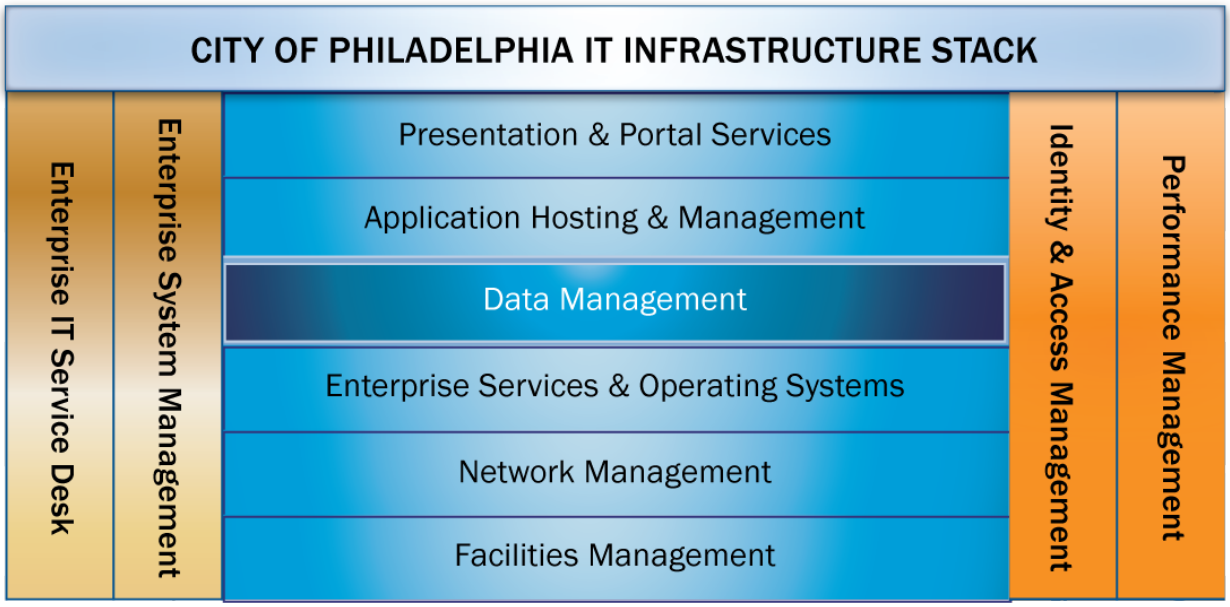

City of Philadelphia Data Architecture

Data Architecture standardizes the design, definition, and relationships of the City's data elements, provides for the governance of those data elements, and guides the creation, maintenance, and availability of the data. Its goal is to make data reusable to the greatest extent possible while improving overall data quality.

Data quality is the common driver for all of the NJSDI components. A primary objective is to first identify the quality of the data within the organization, and then systematically work to improve it.

Data architecture interacts with multiple touch points within the infrastructure, as described below.

Data Modeling

This captures logical and physical definitions of data objects, providing for well-defined non-redundant logical structures that form the basis of all physical database implementations.

Data Collection

This is provided by application development, through acquisition of commercial-off-the-shelf (COTS) software, and by importation of data from external partners and systems.

Data Storage

This manages the life cycle of the data asset at rest. It includes tiered capabilities to meet the storage requirements of different categories of data. It also includes backup, recovery, and restoration capabilities.

Data Transport

This manages the delivery and receipt of data in motion. This can be between internal systems or with external partners. It can use direct writes, pipes, physical media transport, and file transfer protocols.

Data Management (cont.)

Data Integration

This brings together and rationalizes data from two or more systems to create an enhanced data asset not otherwise provided by any one system. It consists of horizontal integration, vertical integration, or both in combination. Horizontal integration is where attributes about an entity in one system are added to different attributes about the same entity in a different system to create a more complete picture (such as appending an employee's payroll attributes to those from human resources). Vertical integration is where additional records of an entity are added to different records about the same entity from a different system to create a larger list of records (such as merging business records from multiple agencies).

Data Publication

This is the delivery of information to different user communities based upon their individual requirements, using graphical end-user tools. The data is formatted as much as possible to anticipate reporting needs, and may be presented differently to different groups, but always from a common *source for consistency.*

Data Governance

Data governance is a set of processes that ensures that important data assets are formally managed throughout an enterprise. Data governance ensures that data is defined, has a known level of quality, and can be used for the intended purpose; in other words, it can be trusted.

City of Philadelphia data governance is focused on identifying those individuals and organizations with the role of defining data objects, identifying the authoritative source for each data object, and classifying each data object. It assists in the resolution of data quality issues, so that City of Philadelphia state government can become more efficient.

Data Steward

The Data Steward is the individual or unit that manages the authoritative source for a particular piece of data and controls its definition and access. A Data Steward is not the same as a Data Custodian, an individual or unit that has been assigned the duty to manage the data under the direction of the Data Steward. A Data Steward is not the same as a Data Owner, which can be a third-party person or organization that the data describes and that has provided the data to the City when requested or required by a State agency.

Data Tiers

City of Philadelphia categorizes data into four tiers – Universal, Enterprise, Line-of-Business, and Programmatic. These data tiers provide a way of framing data governance and data stewardship as well as helping to define the scope of data modeling and data management efforts.

Universal (Tier 0) refers to data commonly referred to as Master Data. This is data that describes persons, places, or things independent of their relationship with the City.

Enterprise (Tier 1) refers to data that is common across all City agencies but within the context of their own organization, such as Financial, Asset, and Human Resources data.

Line-of-Business (Tier 2) refers to data that is common across a particular line-of-business involving more than one agency, such as social services data, business community data, or early childhood data.

Data Management (cont.)
Data Management (cont.)

Programmatic (Tier 3) refers to data that is specific to a single program area within a single agency and is unlikely to have value outside of that context.

Information Classification

Information shall be classified based upon: federal and state laws requiring or prohibiting disclosure; sensitivity, type and importance of the information; business priorities of the agency and the City; and the degree of protection expected for the information. Information shall be classified as confidential, for official use only, or public.

Confidential Information

Confidential information is information that is prohibited from disclosure by law or regulation or by policy or procedure of the City or City agencies, or that is protected from disclosure by the attorney-client privilege or other privilege. Information that is classified as confidential requires the highest level of protection and may be disclosed only to officers, employees, contractors and third party users of the City whose access to the information is strictly necessary to perform their job functions or contracts. Confidential information may not be disclosed to persons or entities that are not officers, employees, contractors or third party users of the City without the express permission of the information owner, unless disclosure is required by law. The information owner is responsible for classifying information as confidential information and for determining the persons and entities that may have access to the information.

Examples of confidential information include, but are not limited to, information covered by nondisclosure agreements or non-disclosure provisions of City contracts, detailed network architecture diagrams, personal information about employees (such as home addresses, phone numbers, and personal family information); information from any employee's personnel file; protected health information (PHI) covered by the Health Insurance Portability and Accountability Act (HIPAA); information relating to legal proceedings involving the City that has not been made public; communication between a City employee and any attorney in the Law Department; information relating to any competitive City procurement of goods or services which, if disclosed, would give an unfair competitive advantage to one vendor over another; information classified as confidential by City-related agencies, where the City information owner determines that confidentiality is appropriate; trade secrets, commercial, financial, or other information of third parties protected by contract or by law; anyone's social security number; security related information for which unauthorized use or disclosure could result in the destruction of life or property and/or disrupt or compromise the security and/or degrade the performance of City networks or information systems; City financial information and reports prior to being publicly released; and all information that is classified as confidential by the information owner.

For Official Use Only

For official use only information is information that is subject to a lesser level of protection than confidential information but that can be made available only to City officers, employees and contractors for which the information is relevant, to the performance of their job functions or contracts. For official use only information may not be disclosed to persons or entities that are not City officers, employees, contractors or third party users without the express permission of the information owner, unless

disclosure is required by law. The information owner is responsible for classifying information as for official use only and for determining the persons and entities that may have access to the information.

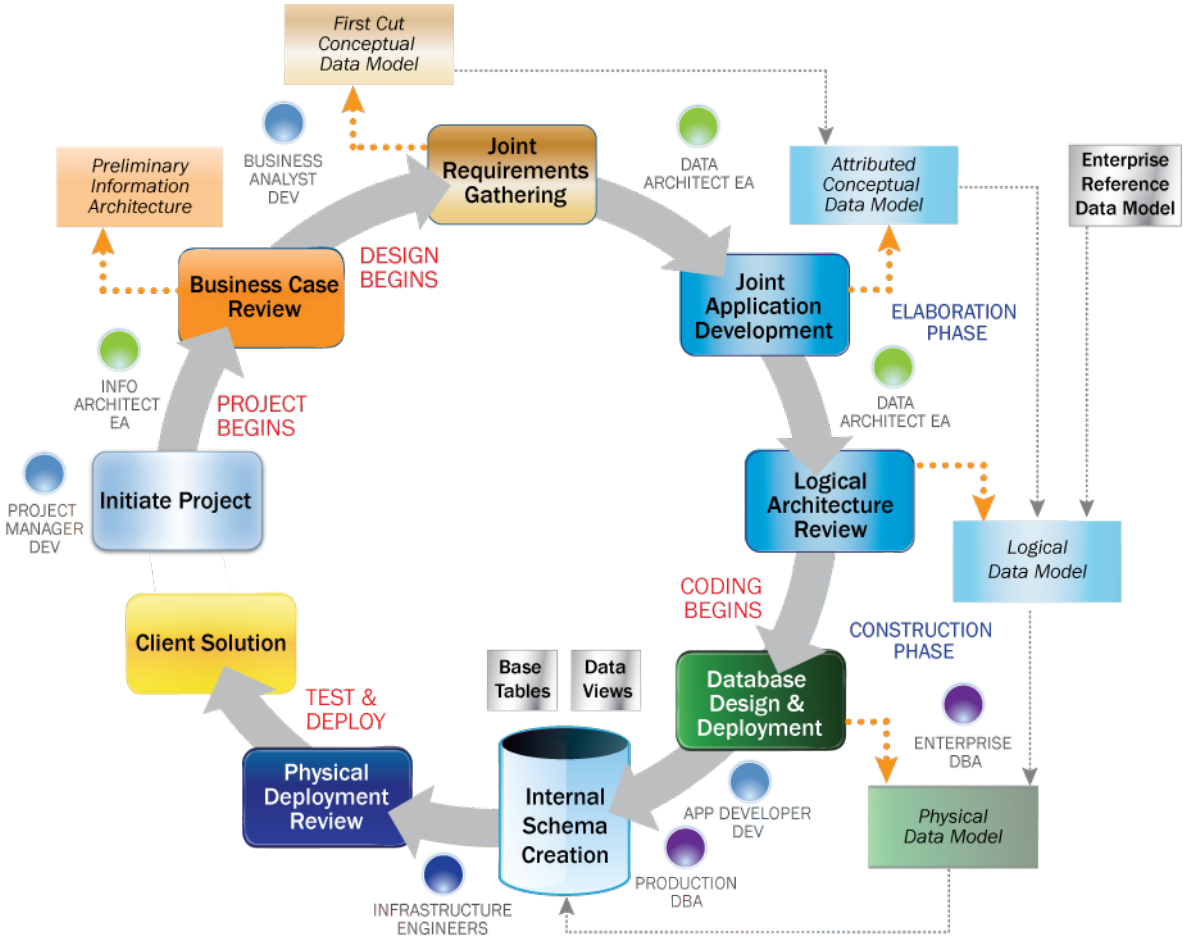

Examples of for official use only information may include, but are not limited to, information on the internal operations of City agencies, internal memos, minutes of meetings, internal project reports, and other documents that are classified as for official use only by the information owner.

Public Information

Subject to subsection (d), Default Classification, information that is not confidential or for official use only shall be classified as public information. Information may be classified as public information, regardless of form or format. The information owner is responsible for classifying information as public information and is responsible for imposing security controls to prevent unauthorized modification and/or

destruction of the information. Examples of public information include, but are not limited to, press releases, brochures, pamphlets, public web pages and materials created for public use.

Default Classification

If information is not expressly classified by the information owner, information shall be treated as for official use only information, whether labeled or not, until it is reviewed and a different classification is applied. Information owners shall ensure that information is classified in a timely manner to prevent the risk of unauthorized access and/or unauthorized disclosure in contravention of its appropriate classification.

City of Philadelphia Information Architecture Design Patterns

A design pattern provides a formal definition of a solution and of the problems to which it applies. The goal of design patterns is to avoid approaching each situation as a problem that has never been seen before and, instead, to make it possible to repeat solutions that have worked. In particular, a design pattern distills the best practices of a community so everyone can apply that expertise. While the approach originated in building architecture and has seen great success in software engineering, design patterns apply equally well to information architecture.

City of Philadelphia has identified these design patterns for different types of information systems.

Transactional System to Collect Data

To the greatest extent possible, new transactional system physical designs shall be developed using a fully normalized logical data model consistent with the City's naming convention. These systems shall be hosted within an industry-standard SQL-enabled relational database management system (RDBMS), and shall use to the greatest extent possible the referential integrity and domain constraint capabilities of the RDBMS to enforce business rules. These systems shall subscribe or consume common reference and master data defined and provided at the enterprise level.

Batch Integration of Inbound Data

Previous assumptions that batch processing windows will always be available to handle any size batch processing requirements are no longer valid. New batch processes must determine if processing smaller batches more often (even in near real-time as batches of one), processing batches while the systems are online, partitioning data or systems, or creating parallel processes are appropriate to achieve the goal of the process.

Real-time Integration of Data

Where there is a need for real-time integration of data, it shall be implemented as a web service. The format for real-time integration shall be defined in XML consistent with the NJERDM. Where one exists, the enterprise service bus (ESB) shall be used.

Provide Data to External Systems from Mainframe Systems

Because data used by one system may be of value to others, and because of the costs associated with creating multiple interfaces on mainframe systems, and because of the complexity of managing outbound interfaces in a mainframe environment, point-to-point solutions shall not be created. Instead, data required by an external system that is not already in the Enterprise Data Warehouse (EDW) environment

Data Management (cont.)

shall be output to the EDW. The external system will either pull or have pushed to it the data from the EDW.

Internal Reporting of Operational Data

Complex reporting needs should not be processed in real-time against critical or already burdened transactional systems. Database tuning for reports is substantially different than for inserts, updates, and deletes (transactions). The type of queries, the volume of the data, and the number of users all add to the processing complexity. Ultimately and invariably, design decisions are made that compromise transaction processing, report processing, or both. Complex reporting must be off-loaded from transactional systems. Techniques include straight replication, the creation of operational reporting marts, and the integration of transactional data into an operational data store. If the same data has a requirement for historical analysis, then the enterprise data warehouse shall be used.

Analytical Reporting against Historic Data

When historical data (defined as the history of changes to a data record, not the history of transactions attached to a current record) is required for analysis, it shall be provided through the enterprise data warehouse environment.

Other types of data exist in the form of snapshots (data that reflects a moment in time, such as a balance sheet), and versions (data that represents the different versions of a record, such as an employee). These data formats are typically not managed in transactional systems. City of Philadelphia manages this data in the enterprise data warehouse in the form of slowly changing dimensions, snapshot fact tables, and profiles. This provides the historical context for reference data.

City of Philadelphia Data Stores

A data store is any database or data repository. Different data stores serve different purposes, and the purpose is independent of the database or repository technology employed. The following specialized data stores are consistent with City IA design patterns.

Transactional Processing Source Systems

These data stores are where the results of business transactions with the City or events of interest to the City are stored. They can be in relational, hierarchical, or file-based database management systems. They can be on a mainframe or on a distributed (network) server. They can be batch processing systems, on-line transactional processing (OLTP) systems, or a hybrid.

Operational Data Store (ODS)

An ODS is a central repository of current operational data initially gathered from a variety of existing transactional systems to present a single rational view of operational data for a single subject area or business unit, or for an entire agency or line-of-business group. History should not be managed or stored in the ODS. Some reporting can occur directly against an ODS, but data can also be replicated into operational reporting areas called Operational Data Marts (Opera Marts).

Data Management (cont.)

Philadelphia Enterprise Data Warehouse (PEDW)

The PEDW is a central repository of historical data that is gathered from a variety of sources to support data integration efforts. An EDW publishes the single version of the truth that supplies historical data to data reusability partners, as well as to analysis areas called Data Marts. It is not a single database, but a consistent data integration environment that consists of multiple subject areas, staging, archiving and persistent storage and multiple physical databases. It is rarely accessed directly by end-users.

The City's information architecture does not support the development of independent data marts (directly built from source systems). Instead, data should be persisted in the EDW for future use. Data is stored in the EDW in one of several ways: in the form of a fully normalized data model for the subject area, as a persistent file en route to a reporting area, as a historical dimension table (reference table with history), as a snapshot table (event table with history), or as a detailed or summarized fact table (array of measure created from the transactional data). Our EDW environment accommodates data for individual subject areas, agencies, and the City as a whole.

Data Mart

A data mart consistent with the City of Philadelphia IA is a pre-defined and pre-formatted subset of data sourced from the EDW or an Operational Data Store that has been identified based on the questions that need to be answered by the report community. Data marts are built for the needs of a specific report community, so the same data may exist in many ways and many combinations in different data marts. They may be logical, consisting of views of enterprise data warehouse data, or physical, consisting of extracts of enterprise data warehouse data. Data is represented in a data mart in one of several ways: in the form provided by the transactional system, as a historical dimension table (reference table with history), as a snapshot table (event table with history), or as a detailed or summarized fact table (an array of measures created from the transactional data).

Dependent data marts always receive data from a consistent, integrated source – never directly from individual operational systems – so the answer to the same question from any data mart is always the same.

City Standard and Supported Technologies

Business Intelligence Publishing Tools

These are query and reporting tools that provide rapid development of reports and can be produced by most business people due to a friendly, graphical interface and a semantic layer than hides the complexity of data relationships from report consumers.

The City does not have a single, standard Business Intelligence Publishing Platform.

Extract, Transform and Load (ETL) Tools

ETL tools are used to move and transform thousands of records in a bulk fashion and are designed and administered in a graphical environment. These tools learn about data and systems and enable reuse of knowledge on subsequent projects.

The City's does not have a standard ETL tool.

Data Management (cont.)

Enterprise Application Integration (EAI) Tools

EAI tools are used to integrate common data across multiple systems at the transaction level, reusing information quality data (metadata). The State requires XML-based web services in a services-oriented architecture (SOA) framework for transaction-level integration.

The City's does not have a standard EAI platform.

Metadata Management

The City of Philadelphia IA requires management of metadata, or information resource data, which can include such diverse categories as data dictionaries, data models, process rules, data lineage, system documentation, transformation rules and security information. Metadata management tools share definitions of data between each other and the systems that they document. When possible, common data names and definitions are shared between systems.

The City's does not have a standard data warehouse metadata management tool.

Data Modeling

Data modeling tools are used to document, locate and reuse data as well as to describe the relationships between data and systems.

The City uses a number of data modeling tools, such as CA ERWin, IBM Rational Architect, Oracle Designer, and Sybase PowerDesigner.

Data Profiling

Data profiling tools are used to discover, document and analyze legacy data, capture metadata, map transformations, and describe the relationships between data and systems.

The City's does not have a standard data-profiling platform.

Data Quality and Cleansing Tools

These tools are used to analyze data values, ensure that data elements are captured and stored in a way to best comply with their business rules and intended application, find patterns of poor quality, standardize addresses, add geographic coding information to records, and perform sophisticated matching of free-form data to find exact or like matches.

The City's does not have any standard data quality and cleansing tools.

Data Mining

Data mining is a sophisticated statistical analysis of data for patterns and clusters. It is not the ability to perform ad hoc queries against data, which is provided by business intelligence tools. Data mining tools can learn from earlier analyses and can look for patterns without guidance.

The City does not have a data-mining standard.

Data Management (cont.)

Supported Database Management Systems (DBMS) Platforms

The State requires that all new transactional database development be built in ANSI SQL-compatible relational database management systems (RDBMS).

The strategic RDBMS products for the City are Oracle and Microsoft SQL Server. While the City is researching open source products such as MySQL and PostGresSQL, they are not part of the City's strategic direction at this time.

The City maintains the following mainframe legacy databases: IBM DB2, IBM IMS, and Software AG Adabas. The City does not anticipate significant new development taking place on any of these platforms, and is engaged in various initiatives to phase out these environments.

Application Hosting & Management

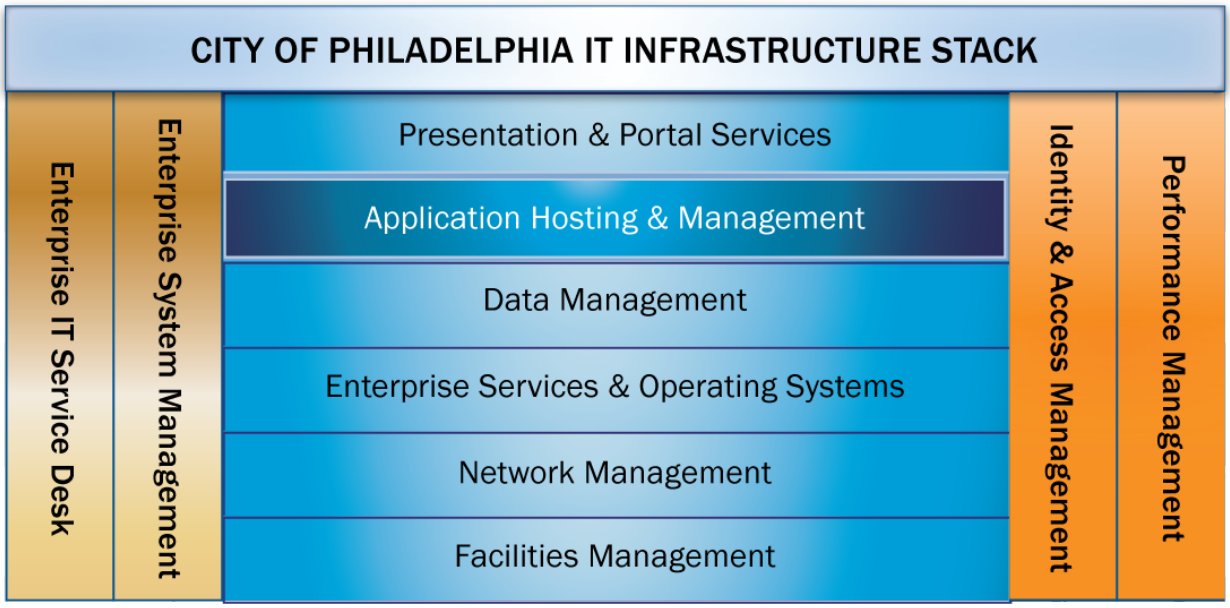

The strategic environments for new applications are service-oriented designs using Microsoft .Net components running on Dell Intel platform servers. All programs should be designed with the goal of developing reusable components. The benefits of building reusable components are evolving into an enterprise framework where common functionality can be shared across applications and platforms. Authentication and authorization should be designed using the City Identity and Access Management Infrastructure.

.Net Application Hosting Environment

The Microsoft .Net application hosting environment uses a multi-tier architecture implementing a web services approach using C# and Visual Basic.NET.

Core Functionality

- .Net framework 2.0, 3.5, and 4.0, which contains Common Language Runtime (CLR) and a collection of .Net application classes
- Internet Information Server 7 (IIS7) is used to host web applications and web services
- Database connectivity to Oracle and Microsoft SQL Server, ODBC, OLE DB and XML data sources
- Industry standard authentication protocols
- Industry standard cryptography features for encryption, digital signatures, hashing and random number generation

High Availability

- Separate Business Logic and Persistence (Data) Tiers. This enables greater scalability across both the business logic and persistence tier while allowing for integrated installation and administration
- Drive Redundancy. Each server contains two mirrored drives and a hot spare which allows the server to continue functioning even if two drives are lost

Application Hosting & Management (cont.)

- Server Redundancy. There are duplicate servers in both the public and secure tiers to enable workload balancing and continuous availability in the event of a server failure
- Horizontal Scalability. As the workload increases, additional servers for application and web support can be easily added
- Network Load Balancing
- SSL Offloading
- Tivoli tools are used to monitor the health of servers to detect and correct problems

Document Management

The City of Philadelphia (the City) has in place resources and operations for the processing and management of electronic documents.

Automated document management/storage systems include, but are not limited to, systems based on electronic workflow automation, on-line storage and retrieval of record images, Internet-based filing/record retrieval, or combinations using technological platforms such as these. City Enterprise Services include mail processing, remittance processing, document screening/preparation, electronic scanning, index/application data capture, and hosting of electronic images on server platforms.

In virtually all-new systems there are potential elements for document management functions. Agencies should seek to utilize existing City document management services as a first choice rather than acquiring or building duplicative services models.

Records Retention Schedules and Requirements

Proposed systems should provide for and adhere to the City's retention schedule and information classification requirements. Contact the Records Department for further details regarding the City's records retention policies.

Technology Infrastructure

While the City may have various implementation of vendor software, which accomplish scanning and imaging operations, the primary software that is in use for document imaging is the Documentum product line. Application integration for scanning and imaging solutions will utilize interfaces into the Documentum software where they are to utilize existing services.

Legacy and Mainframe Services

The City has IBM enterprise servers which host applications for the law enforcement community, tax systems, and payroll among others. The mainframes are geared toward high volume activity and have excellent response time and availability track records. Geographic Information System (GIS) Services

The City has a goal of spatially enabling any application that would benefit from geo-awareness. The City definition of spatially enabled means that the system is capable of:

- integrating spatial data (e.g., data with a location component) with other business data across multiple, heterogeneous data sources
- supporting abstract data types (e.g., images, text, and spatial data), spatial operators and functions, and spatial locator indexes.

Application Hosting & Management (cont.)

Managing and accessing spatial data across the City's IT enterprise is facilitated through a gateway which utilizes a combination of technologies including Environmental Systems Research Institute (ESRI) Arc Spatial Data Engine (ArcSDE). Spatial data is served up in a format that can be accessed by a variety of desktop GIS clients, served out to the Internet using ESRI's ArcGIS Server technology or by other applications using standard SQL queries. Spatial data is hosted on an Oracle platform providing reliability and scalability.

ArcGIS Server technology provides the foundation for distributing high-end geographic information systems (GIS) and mapping services via the Internet. This technology also enables users to integrate local data sources with Internet data sources for display, query, and analysis in a Web browser. We utilize ESRI's ArcGIS Server, a powerful, scalable, standards-based tool used to design and manage web services for map display and geoprocessing. This technology is currently integrated in the City's Shared Server Infrastructure (SSI) using a three-tier application architecture. ArcGIS Server is maintained at a release level at or near the latest available.

For ArcGIS Server application development, all APIs provided by ArcGIS Server are available in the City's environment, but among the several REST API clients, JavaScript is preferred. Silverlight or Flex are available if necessary, but are discouraged because of the limited client platforms supported. The city does not support the use of the ESRI Web ADF in any new applications as this technology has been deprecated. The city also prefers the use of the ESRI REST API over the ESRI SOAP API for data access.

An array of web services is maintained on the shared infrastructure to meet common functional requirements such as address geocoding, reprojection of data between different coordinate systems, and other similar tasks.

Any proposed solution that includes a GIS component and/or incorporates spatial data is evaluated, planned, designed, and implemented in concert with the OIT GIS Services Group. Applications that are geo-enabled are in compliance with the OpenGIS Consortium specifications for spatial data (http://www.opengis.org/). The City's preferred GIS software platform is the ESRI set of products and tools (http://www.esri.com/).

ePay Gateway

OIT maintains a set of enterprise ePayment web components and services that provide Internet based payment processing to City agency applications. The ePay Gateway allows custom developed Web based applications to either process:

- Credit card transactions by interfacing with a payment gateway provider; or
- eCheck transactions by allowing governmental entities to accept electronic checks via the Internet

Implementation of the ePay Gateway is facilitated through web protocols. As such, they can be used with any compliant application in the .NET and J2EE environments.

Single Sign-On

See section on Identity Management, Authentication & Authorization Services.

Presentation & Portal Services

Web Servers

Anonymous access to the City's static public information is provided through the public access Web servers (www.phila.gov).

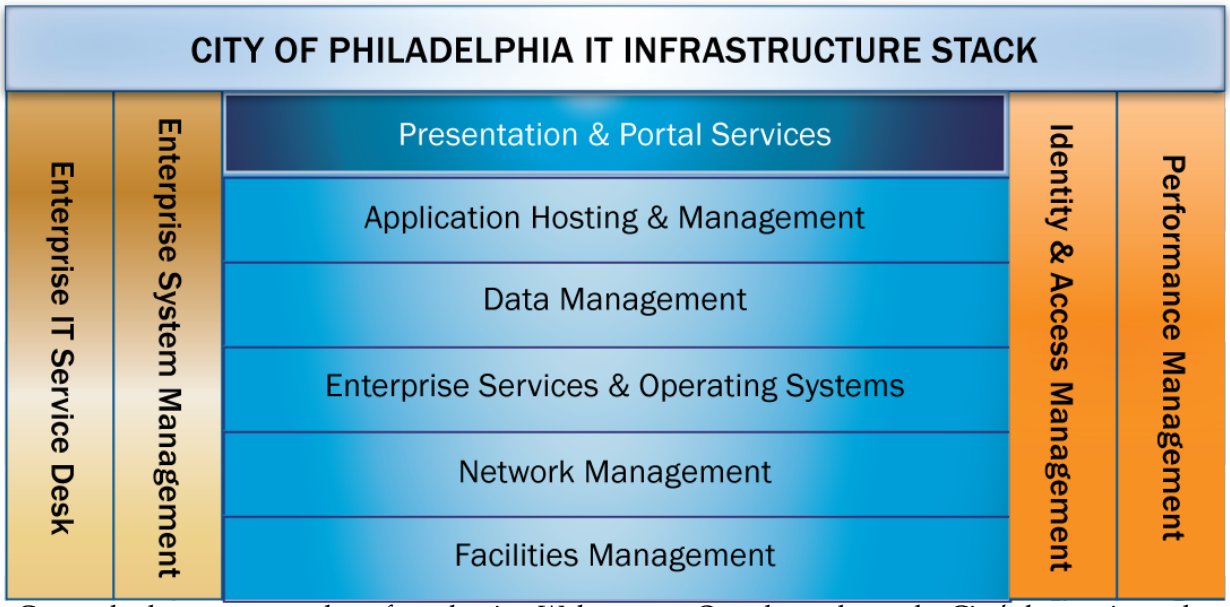

Currently there are a number of production Web servers. One cluster hosts the City's home site and related flat file information (www.phila.gov) as well as several ASP.NET web applications. One cluster supports ASP.NET for serving secure applications. A SharePoint farm supports web content for City agencies that leverage the City's web content management services.

The web server platform is Internet Information Services server. It provides the following capabilities to City agency developers:

Web Application Development

- Support for ASP.NET (2.0+), Javascript, HTML5, and CSS3
- Session management service to track information for specific users
- Reuse of applications and components that are developed separately
- Server-side preprocessing of content using SHTML

Reliability and Availability

- High server uptime through multi-processing mode and process monitors
- Load balancing configuration with Big-IP F5 Local Traffic Manager for high availability

Management and Administration

- Integration with City's IAM
- Command-line interface for HTTP server administration, certificate and key management, and Web application deployment

Presentation & Portal Services (cont.)

Performance and Scalability

- High performance through an advanced multiprocessing, multithreaded architecture; efficient use of kernel threads; and sophisticated memory management, Server-side HTML (SHTML) and chunked encoding to enhance the performance of dynamic content
- HTTP 1.1 and HTTP compression
- Scalable, keep-alive handling

Web Content Management

SharePoint Server provides enterprise web content management services to City agencies. SharePoint allows web designers to control the look and feel of the finished pages while allowing non-technical users to provide the content that appears in the final product. Application Infrastructure Services provides the technical support for the infrastructure, and Creative Services provides the development expertise by creating the page templates and is responsible for end user training.

Enterprise IT Service Desk

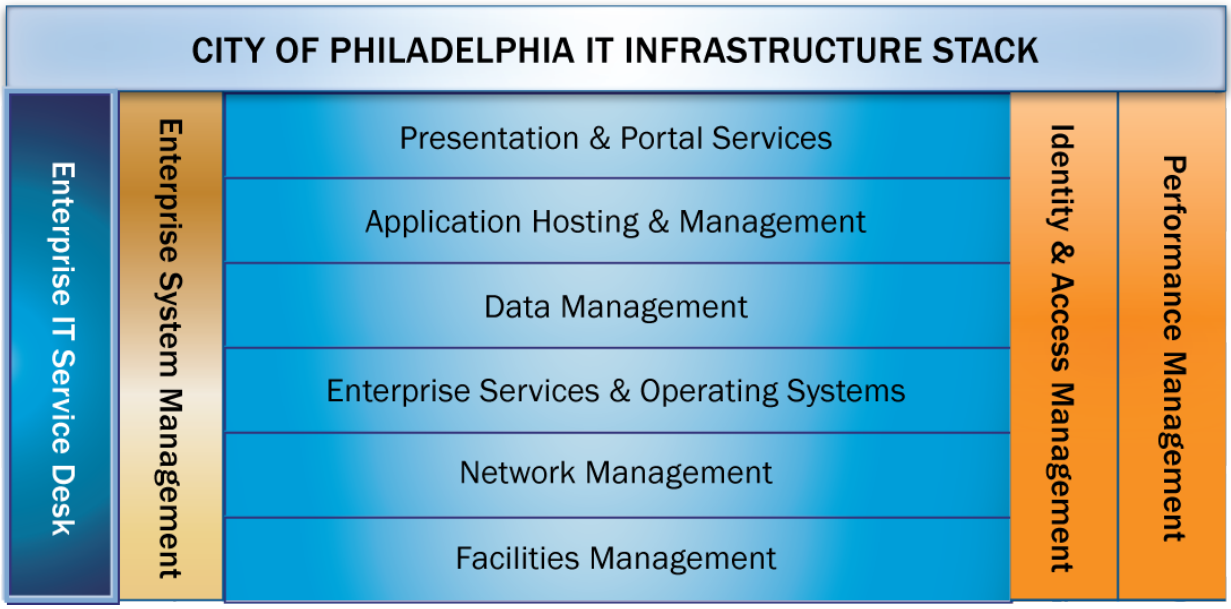

The Enterprise IT Service Desk provides City agencies with a single point of contact for technical incidents and advocacy for response and resolution. The Enterprise IT Service Desk is staffed 24x7 to resolve system outages. Request for IT support submitted off hours are forwarded to the Operations Unit staff by dialing 686-8213. All requests made to the Enterprise IT Service Desk are recorded in the service desk request tracking system. The system simultaneously e-mails the resources that have been identified to resolve specific incidents. Resources begin the incident resolution process and update the request with status information until it is resolved. System users can access this system via a web browser to monitor the resolution status of their incident.

Business Hours

The normal business hours are Monday thru Friday from 7:00 AM to 6:00 PM with the exception of City holidays.

Submitting Requests

There are several means for contacting the Enterprise IT Service Desk. Listed below are the various methods along with the Service Level for Service Desk response.

Enterprise IT Service Desk (cont.)

Method	Details	Response Time
Online (Intranet only) http://help.city.phila.local/	Here requests can be submitted online. When a request is submitted online a service ticket is automatically generated in the service request tracking system. The Service Desk will assign tickets to the appropriate resources.	Service Request created immediately
By Email DoTHelp@phila.gov	The Enterprise IT Service Desk will create and assign tickets on behalf of the customer.	**During Business Hours** Non-critical: Within 2 hours of receiving an email Critical: Within 1/2 hour **Off Hours** Non-critical: Within 4 hours of receiving an email Critical: Within 2 hours
By Phone (215) 686-8213	The Enterprise IT Service Desk will create and assign tickets on behalf of the customer. Callers may choose to leave a voicemail or hear a system status update for major disruptions or outages.	Calls are answered in less than 90 seconds. Will respond to voicemail within 1 hour

Customer Service Surveys

In order to improve service levels, the Enterprise IT Service Desk sends customer service surveys to customers once a ticket is closed. Customer feedback is used to improve service.

Performance Management

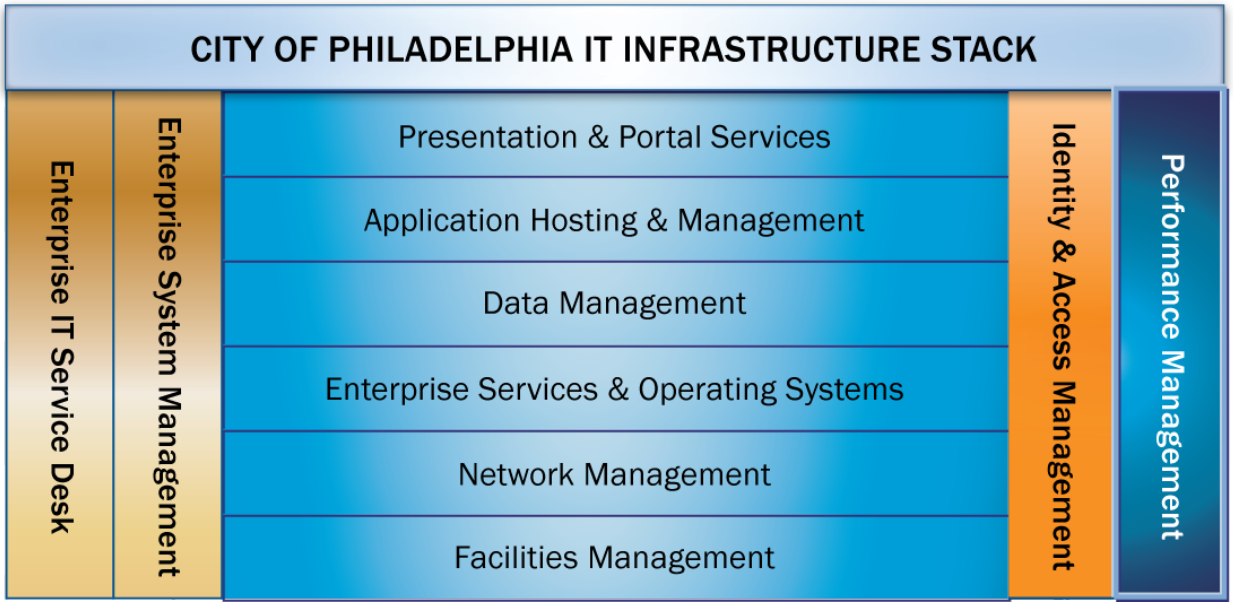

Application Instrumentation and Monitoring

The Office of Innovation & Technology (OIT) is in the process of creating a holistic, enterprise strategy for performance monitoring. This strategy will address monitoring of the following:

- ASP.NET applications, including web services and SharePoint solutions
- Oracle Database Servers
- SQL Server Instances

OIT currently provides support for the Logging and Exception Handling Application Blocks from the Microsoft Enterprise Library 5.0. Leveraging these Application Blocks, ASP.NET applications can be configured to record log entries to a database server, the Windows Event Log, and to send log entries as email messages.

Network Performance

Network Performance, Application Triage and Performance Service Level Monitoring

The OIT currently utilizes multiple product sets to monitor, assess, diagnose and provide fault domain or root cause analysis for network and application performance issues. Software and hardware based network and application probes are deployed to perform baseline analysis of existing network environments prior to deploying new applications, upgrading existing applications or implementing new IP services such as VOIP or telephony applications. Application protocols, their respective traffic volumes traversing the local (LAN) and wide area network (WAN) are identified and their bandwidth consumption, average response times, Round Trip Times (RTT) (TCP handshake), conversation pairs and traffic volumes measured. This analysis can be used as a benchmark comparison against future application or network performance.

Performance Management (cont.)

Application pre-deployment assessment services are available for departmental and agency clients to assess the network performance characteristics of individual application functions or transactions. Characteristics such as TCP windowing, packet size, conversation flow, node sending and processing behaviors and inspection of processing threads between clients and servers or between nodes in a n-tier hosting environment are examined. The assessment process also includes performance modeling based upon variations in available bandwidth, processing performance and modification in TCP/IP stack parameters.

In the production environment, application triage tools are deployed to monitor and diagnose degradations in application performance and to determine the root cause(s) of poor application performance. Triage tools help to determine whether poor application performance are the result of underpowered hosts, the network infrastructure, the application code or an inefficient host server platform/OS or database.

Pro-active/Real-time monitoring of Application Service Levels is the newest offering available to departmental and agency clients who have deployed application systems within the n-tier environment at the State's data centers. Deployed performance Service Level Management (SLM) toolsets monitor end-to-end flow of application traffic traversing the E-commerce environment. The SLM process provides both real-time and historical information regarding the performance usage and availability of key business applications (see Appendix 4 –Service Level Management Toolset). The SLM toolsets provide deep dive analysis capabilities of Web based application components including SSL decryption, HTTP page load sequencing, hit level and error analysis, application protocol analysis and decode, including database performance metrics, tiered fault domain isolation analysis, available in real-time or through an extensive set of data mining services.

Network Monitoring

The performance of network appliances and services is monitored in terms of the utilization of CPU, disk space, interfaces, and memory as well as ping latency.

Vulnerability Management Services

As required by policy and procedures, the OIT utilizes vulnerability management as a measure to keep key resources within the Garden State Network safe from hacking and Internet cyber-attacks. The OIT also oversees vulnerability management efforts in order to ensure City of Philadelphia Government (the City) Executive Branch agencies and agencies are meeting policies, regulations, and directives required by the City, the State of Pennsylvania, the U.S. Federal Government, and private industry. To control and manage risk attributed to security vulnerabilities, the OIT provides an Enterprise Vulnerability Management system to departments and agencies. The system is utilized for testing new hardware introduced into network infrastructure and provides an immediate view of network security and compliance posture. The vulnerability management system is also capable of auditing and assessing networks for the possibility of weaknesses that tend to be channels for data and information theft, unauthorized access, or targeted exploitation. Use of the vulnerability management system is guided by the workflow process of detection, removal, testing, and control.