Architecture

Advanced Design and Setup Guide

Version: Foundation 22.1

Written by: Product Knowledge, R&D
Date: February 2022

Hyland

Copyright

Information in this document is subject to change without notice. The software described in this document is furnished only under a separate license agreement and may be used or copied only according to the terms of such agreement. It is against the law to copy the software except as specifically allowed in the license agreement. This document or accompanying materials contains certain information which is confidential information of Hyland Software, Inc. and its affiliates, and which is subject to the confidentiality provisions agreed to by you.

All data, names, and formats used in this document's examples are fictitious unless noted otherwise. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright law, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Hyland Software, Inc. or one of its affiliates.

Hyland® and Hyland Software®, as well as Hyland product names, are registered and/or unregistered trademarks of Hyland Software, Inc. and its affiliates in the United States and other countries. All other trademarks, service marks, trade names and products of other companies are the property of their respective owners.

© 2022 Hyland Software, Inc. and its affiliates. All rights reserved.

Table of Contents

Copyright **2**
Goals of the Perceptive Content Architecture **5**
Perceptive Content Architecture: a Logical Perspective **5**
The Client Tier 6
Perceptive Content Client *6*
Web services *6*
The Business Tier 6
Document Management *7*
Document Imaging and Capture *7*
Workflow (Business Process Management) *7*
Metadata and Full Text Search *7*
Operational Reporting *7*
The Data Tier 8
OSM *8*
Database *8*
The Design Pillars of the Logical Perspective 8
Perceptive Content Architecture: a Physical Perspective **13**
Client Components 14
Perceptive Content Client *14*
Interact Clients *14*
Scanner *14*
Web service clients *14*
Middle-tier Components 14
Integration Server *14*
Server Components 14
Perceptive Content Server *15*
Database *15*
OSM *15*
Network Information *15*
The Design Pillars of the Physical Perspective 16
Perceptive Content Architecture: a Security Perspective **19**
Content Security 19
Digital Signatures 19

Administering Security in Perceptive Content 20
The Perceptive Manager Role *20*
The Department Manager Role *20*
Establishing Users and Groups *20*
Authentication and Client Validation 20
Encryption 21
Auditing 21

Goals of the Perceptive Content Architecture

The Perceptive Content difference is rooted in its careful architecture. The Perceptive Content product suite provides a carefully tailored feature set developed according to the design pillars of flexibility, functionality, usability, scalability, extensibility, and dependability. By providing a product that combines these key design elements, Perceptive Software is able to deliver customer-centric document management, imaging, and workflow software that is unmatched in functionality, performance, and value.

The fundamental pillars of our software design and development philosophy are:

Flexibility: Provide the installation, deployment, and configuration options necessary for our diverse customers to be successful and deliver the product when customers need it.

Functionality: Provide a feature-rich set of products for a document management, imaging, and workflow software our customers can use today while ensuring that future plans are within reach.

Usability: Enable users to quickly take advantage of the flexibility and functionality of the Perceptive Content system by creating solutions using proven standards and intuitive interface design.

Scalability: Design Perceptive Content to scale to the number of users, types of processes, and volume of data our customers require today while ensuring future needs can be met.

Extensability: Provide technology that makes it simple to extend Perceptive Content functionality through optional software components. Also, enable Perceptive Content to easily reach out and interact with your existing systems through a variety of patented and standard integration choices.

Dependability: Make Perceptive Content robust and secure in all aspects, ensuring high-availability.

The architecture information provided in this document describes how Perceptive Content meets these design pillar goals, based on logical, physical, and security perspectives. The logical section provides information about the Perceptive Content architecture for a target audience of system developers and integration personnel. The physical section contains information about the Perceptive Content components and network for IT, hardware, and network personnel. A security section includes considerations for both physical and application (logical) security. This section is intended for both developers and IT personnel.

Perceptive Content Architecture: a Logical Perspective

Conceptually, the Perceptive Content architecture consists of the following logical tiers:

The Client Tier consists of our out-of-the-box client applications, which provide rich user interfaces for Windows desktop and browser users, and the class of clients that can be written by our customers and partners using web services.

The Business Tier is the core of the Perceptive Content document management, imaging, and workflow system. This tier processes requests from the Client Tier, provides the business logic to perform the requests, and communicates with the Data Tier to search, retrieve, manipulate, and store document images and associated metadata. This tier is wrapped by a rich set of web services that enable customers to easily extend the product and use it as content infrastructure within their larger enterprise systems.

The Data Tier consists of the components that store data: the object repository for the document objects and a relational database for the corresponding metadata.

These logical tiers work together to provide the features, performance, and security you need for a best-of-breed document management, imaging, and workflow system delivered more quickly than any comparable system in the market today.

The Client Tier

Perceptive Content Client

The Perceptive Content Client is the graphical user interface into your document management, imaging, and workflow system. The Perceptive Content Client gives users single-click access to any piece of unstructured data (scanned image, fax, email, enterprise report, or imported file) directly from records in your business systems. The Perceptive Content Client is written in C++ to ensure quality desktop integration.

Web services

Integration Server makes Perceptive Content content and functionality available to third-party applications by enabling external customer applications that are coded in standard development languages that are also compatible with HTTP Web services, to send and receive data from Perceptive Content Server. Integration Server exposes Perceptive Content functionality through a multiplex of resources, such as Document resources, Folder resources, and Drawer resources. Each Web resource includes a family of functions that Integration Server makes available to third-party applications, such as starting a session, copying a document, or routing a document.

The Business Tier

Perceptive Software has carefully designed Perceptive Content to be modular and flexible so that the components of the system run on a single server class machine or distributed to run across multiple server machines, depending on the scale and volume needs of the customer. The Perceptive Content architecture ensures that the Perceptive Content Server manages and executes all client requests from

the Business Tier. There are two benefits to this architecture: first, it eliminates any client-side, database server connectivity requirements, and second, it allows the collocation or distribution of all key server applications, which can run on and communicate with different servers and operating systems.

The Business Tier provides a rich set of functionality to the Client Tier as described in the following paragraphs.

Document Management

Document Management services provide you with single-click access to relevant content at any point within your business process. Additional document management features include library services that enable users to create, edit, and maintain a document in accordance with your normal business process workflow. Your users can easily version control a document by checking it out. Then, they can make changes to the document and check the document back into the system when they finish editing it. These services ensure the integrity of the document, while managing the changes and ensuring all compliance initiatives. A complete audit trail and history is available for all documents in the system.

Document Imaging and Capture

Capture services provide the means to capture any kind of content— electronic documents, document images, content assets such as graphics files, and extracted content from structured and unstructured documents, from any source (scanners, faxes, emails, file system, print drivers, and programmatically), at any time, and at any volume, and then automatically deliver that information in full or by its extracted parts to the wider systems that need it.

Key components of the capture system are CaptureNow, which provides the user controls to specify how to capture documents, and various recognition technologies to OCR, ICR, and OMR documents for purposes of automating, indexing, content acquisition for full-text searching, and extracting line-item level detail from structured and unstructured documents. In addition, robust capture services are provided for obtaining enterprise report data and EDI files.

Workflow (Business Process Management)

Workflow services provide a powerful business process modeling and execution environment that enables you to place content into workflow and route it through your business processes according to business rules defined using a combination of natural language rules and custom scripting all within simple and elegant workflow design studio. As a workflow user, you can add a document to workflow, route it, respond to alarms and messages about the status of your items, and process it by, for example, viewing, annotating, and digitally signing it.

Metadata and Full Text Search

Retrieval services enable you to find a document that meets a certain criteria by searching through index keys and searching through Enterprise Reports Management (ERM) object index keys.

Operational Reporting

Perceptive Content Business Insight allows you to present your Perceptive Content data using a ready-to-run library of reports as well as the reports created by your Business Insight Author user. In Business Insight, you can view, run, schedule, and archive instances of the reports used in your enterprise. Flexible output allows you to run, view, and save instances of a report as an HTML page, a PDF document, a comma-separated file (CSV), an XML file, or as a Microsoft Excel spreadsheet. You can then distribute instances of the report instantly using email, shared file locations, or FTP.

With Business Insight, high-level dashboards, self-service report tools, and other flexible options enable your organization to evaluate and improve the efficiency of content-related routines. Perceptive Content Business Insight allows you to present your Perceptive Content data using a ready-to-run library of reports as well as the reports created by your Business Insight Author user.

The Data Tier

The Data Tier is comprised of an extensible repository to store files from various content, including scanned images, electronic documents, and forms. This repository is a hybrid storage model comprised of file-system storage, the Object Storage Manager (OSM), which stores the content and a relational database that stores the metadata related to the content.

OSM

The OSM can store unlimited amounts of content in its original format, for example, TIFF, PDF, or Microsoft Word. By storing every page of a document as a discrete object, the OSM enables Perceptive Content to deliver pages to users on demand. The OSM is a tree structure file system that consists of a main directory comprised of sets or branches. Each OSM set can contain 134,217,728 documents or pieces of content. As your storage needs increase, you can configure Perceptive Content to add additional OSM sets automatically.

You can configure the OSM to store objects across any number of file systems on a variety of platforms and architectural designs. This efficient, hybrid storage model ensures the Perceptive Content database maintains steady performance even as usage and database size grows.

Database

The Perceptive Content database stores the metadata information related to the content stored in the OSM, as well as system information. The Perceptive Content database uses the optimal number of indexes necessary to increase read performance without negatively impacting record inserts. Thorough index auditing supports this design, where the Perceptive Content database leverages the unique database tuning tools of the chosen platform.

The Perceptive Content Server performs all database transactions, a factor that contributes to scalability. The Perceptive Content Server creates a pool of database connections, which eliminates any client-side database connectivity requirements. By establishing a connection pool to the database on behalf of the clients, Perceptive Content maximizes system resources on the server, which enables it to handle increasing user loads while ensuring linear performance.

All internal database communication occurs through industry-standard SQL, instead of DBMS-specific code used in stored procedures. This design streamlines the code for all DBMSs and increases performance by taking advantage of the power and flexibility of an integrated programming language versus a database-specific interface language. An additional advantage this design provides is that no maintenance is required to port stored procedures when manufacturers release new database versions.

The Design Pillars of the Logical Perspective

The following table summarizes the design pillars of Perceptive Content as they relate to the logical perspective of the architecture.

Pillar	Description
Flexibility	**Similar Client Design** The Client Tier products provide flexibility because almost all Perceptive Content functionality is available through a Windows or web browser client, or you can write your own client using web services. For example, you can search, view, and manage content in a Windows environment using Experience Client. **OSM Storage Options** In the Data Tier, the file-system based OSM provides you with the flexibility to store your data using a wide array of commercial technologies that best meet your performance and cost parameters.
Functionality	**Features** The Client Tier gives users single-click access to any piece of content through the Business Tier functionality (scanned image, fax, email, enterprise report, or imported file) directly from records in your business systems. The Client Tier provides access to all functionality, including the capacity to manage and annotate documents, capture, view, search, route objects in workflows, and provide reports.
Usability	**Familiar Platforms** Perceptive Software built the Client Tier on platforms that are familiar to our customers, including Windows, web browser-based, and web services-based solutions. Providing the clients in familiar platforms makes them more usable for our customer base. The clients adhere to graphical user interface standards that provide familiar tools, such as windows, menus, and dialog boxes. **Consistency** The Perceptive Content provides a consistent look and feel so that customers can seamlessly move between platforms. In addition, Perceptive Content Client includes standardized, information-rich, help topics that provide assistance when needed.

Pillar	Description
Scalability	**Multi-threaded Approach** In the Business Tier, Perceptive Content uses a multi-threaded server model to handle large volumes of requests while scaling very efficiently. Using this approach, Perceptive Content tracks multiple users and their requests as threads with separate identities. This approach enables Perceptive Content to address an unlimited number of requests at the same time, while tracking the status of work for each thread until the work finishes. **Hybrid Storage Model** In the database, Perceptive Content stores metadata (information about the document such as, document indexes, version information, security privileges, and workflow status) and unlimited amounts of customer defined data, while storing document objects on the file-system rather than binary large objects (BLOBs) in the database. This hybrid storage model significantly lowers the database size, and increases overall system performance. **Flexible Storage** In the Data Tier, for the OSM, you can create separate object storage structures that you can subdivide and cluster using virtually any logical construct. For example, you can optimize storage around patterns of usage by creating multiple storage locations for different document sets. Using multiple storage locations, you can place frequently accessed data in one file system, while placing other data in near-line storage. This division allows the Perceptive Content object store to be distributed to extremely high-performance SANs and other dedicated I/O maximized devices, such as RAID, optical disks, HSM systems, or any combination of storage devices. Perceptive Content also supports the ability to archive objects to CAS (Content-Addressable Storage) devices. **Remote Services** Over TCP/IP, any Perceptive Content Client anywhere in the world can securely communicate with any Perceptive Content Server using your LAN, WAN, intranet, extranet or the Internet itself. The Perceptive Content communications subsystem bursts packets efficiently between the Perceptive Content Client pool and the Perceptive Content Server, optimizing load handling and providing an excellent user experience with sub-second content retrieval. Because Perceptive Content is a pure TCP/IP-based product, it offers the highest level of raw network performance and the greatest breadth of back-end scalability.

Pillar	Description
Extensibility	**Integration in the Client Tier** The Client Tier offers a powerful variety of non-programmatic integration options collectively known as LearnMode that set it apart from stand-alone, post-processed, “filing cabinet” solutions. The Client Tier uses LearnMode for integrating with virtually any application delivered to any desktop application or enterprise system. Client-side integration options support single-click indexing and retrieval as well as user collaboration via email. Perceptive Content uses MAPI to invoke whatever email client you use and to auto-populate the body of the email. All major Perceptive Content functions are accessible programmatically on the desktop via Visual Basic script and COM function calls to allow a variety of more advanced integration options. **Integration in the Business Tier** The Business Tier provides functionality for customers who want to programmatically integrate Perceptive Content with their systems. Integration Server can be used with third-party applications to flexibly incorporate Perceptive Content functionality directly into products and enterprise systems by providing the required classes, functions, methods, subroutines, or other routines. Integration Server interacts seamlessly with Perceptive Content Server to increase efficiency with the use of your existing development tools. iScript enables the customization of business processes in Perceptive Content and provides the means for those processes to interact with third party systems using industry-standard interfaces such as COM, ODBC, ADO, and SMTP.

Pillar	Description
Dependability	**Proven Programming Languages and Standards** Perceptive Software designed and built Perceptive Content using C++ and Java 2 technologies following object-oriented principles to provide maximum performance. Perceptive Software designs products that adhere to industry standards, such as: • ODBC (Open Database Connectivity), the industry-standard method of sharing data between the database and the Business Tier in Perceptive Content. • XML (Extensible Markup Language) and XSL (Extensible Style Language), the W3C (World Wide Web Consortium) standards you use in Perceptive Content when you create worksheets that design content through customized tags and styles. • SOAP (Simple Object Access Protocol), a standardized communication protocol, enables programs running under different operating systems in a network to communicate with each other. • REST (Representational State Transfer), an architectural style for message exchange that addresses the Web as remote resources. In a RESTful application such as Integration Server, each URL points to a resource. • WSDL (Web Services Description Language), an XML-based language used to describe and access services electronically. • iScript, an ECMA-262-compliant JavaScript based language. **OSM Mirroring** In the Data Tier, the OSM provides the ability to write object data to multiple sources simultaneously to provide real-time redundancy of the OSM, which is also known as mirroring. You can configure Perceptive Content storage locations on the same server as the Perceptive Content Server or on an external file system located anywhere on your LAN or WAN. For example, one file system can be on RAID 5 with its mirrored file system located on an optical jukebox, while simultaneously writing to a magnetic storage system in an offsite, secure location. If the drive where the primary OSM storage fails after you select OSM mirroring, Perceptive Content automatically attempts to read the data from the secondary location.

Perceptive Content Architecture: a Physical Perspective

Perceptive Content is a true n-tier application. This architecture eliminates any client-side, database server connectivity requirements, allows the collocation or distribution of all key server applications across different operating systems, and minimizes the need for database specific business logic while eliminating the need for stored procedures.

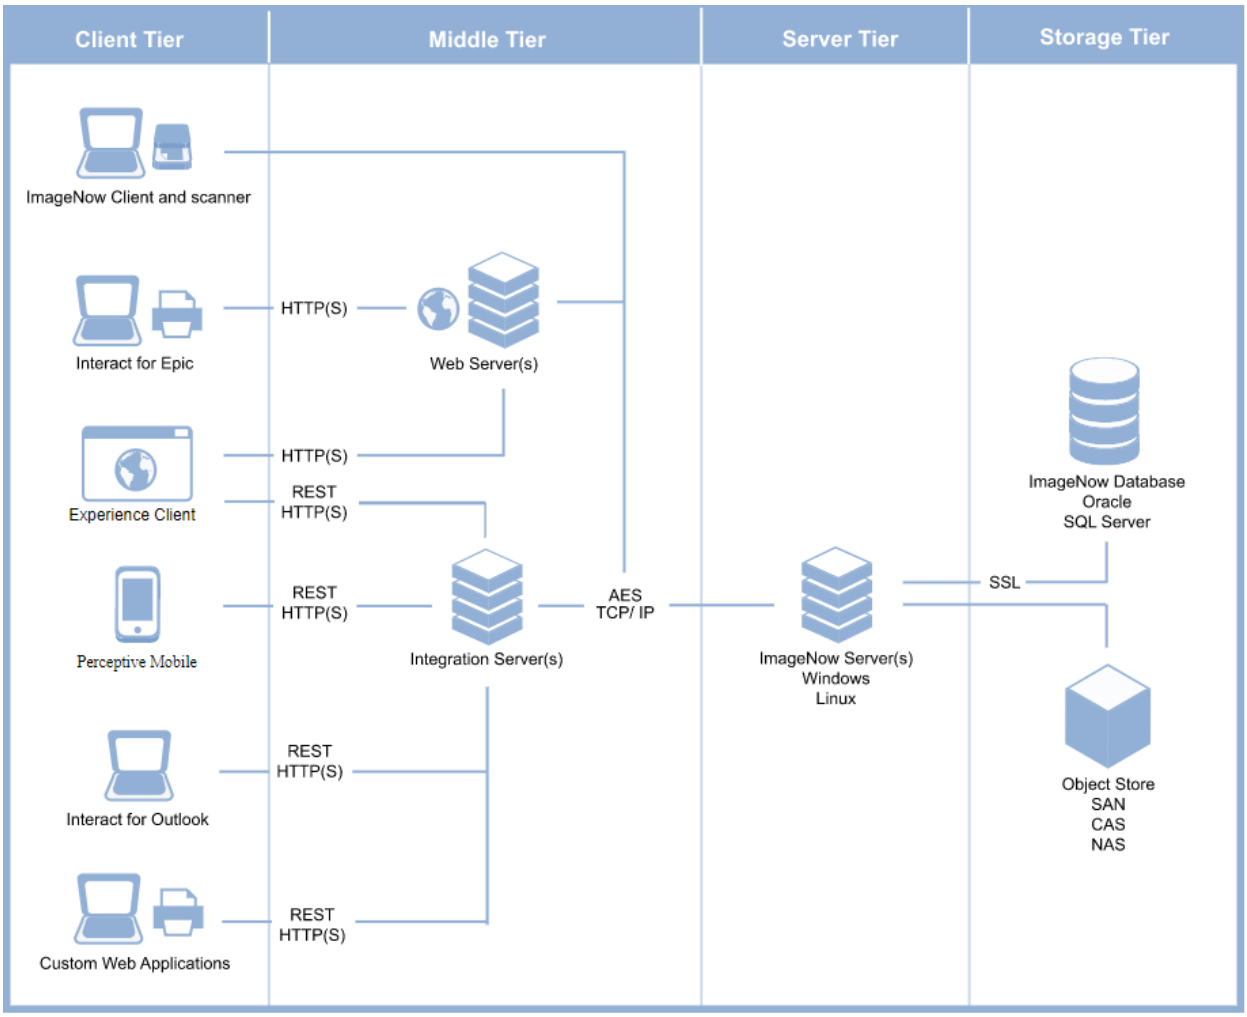

Communication between Perceptive Content Server and other components occurs over TCP/IP using a socket-based communication protocol defined by Perceptive Software to achieve best possible performance. The socket communication consists of a series of calls that the clients use to request information from the Perceptive Content Server.

The socket communication between Perceptive Content Server and Integration Server implements AES (Advanced Encryption Standard), to encrypt the transport channel. The Perceptive Content Server provides the content by accessing the database and the OSM file system.

The following sections describe the physical aspects of the software, as well as hardware and network information. For the most up-to-date hardware information, refer to the *Product Technical Specifications* document.

Client Components

The Client Components include the Perceptive Content Client, Interact Clients and Web service client.

Perceptive Content Client

The Perceptive Content Client is a desktop interface that runs on Windows platforms using standard TCP/IP networking. Perceptive Software writes this client specifically for the Windows environment, using high performance C++. You connect the Perceptive Content Client to the Perceptive Content Server using standard TCP/IP.

Interact Clients

Interact Clients provides embedded integration to your Perceptive Content documents from within business and collaboration applications. Perceptive Software offers several Interact Clients, such as Interact for Epic and Interact for Office. The integration method depends on the type of application. For example, Interact for Epic uses HL7 messages to handle the communication between Epic and the Perceptive Content Server. For product technical specifications and system requirements, refer to *Product Technical Specifications*.

Scanner

The scanning environment in the Perceptive Content Client, uses drivers from both Kofax and EMC and supports a wide range of scanners from Fujitsu, Canon, Bell & Howell and other manufacturers to ensure the most flexible batch and ad hoc capture environment possible. The Perceptive Content Client running a scanner is the "scanstation." The scanstation computer must contain a slot to accept a Kofax Adrenaline or SCSI boards to drive each scanner or be USB-enabled.

Web service clients

Integration Server makes Perceptive Content content and functionality available to third-party applications by enabling external customer applications that are coded in standard development languages, such as Java, C++, or C#, and that are also compatible with HTTP Web services, to send and receive data from Perceptive Content Server.

Middle-tier Components

The Middle-tier Components includes Integration Server.

Integration Server

Integration Server is a middle-tier Web service that provides communication over a network between Perceptive Content software and third-party applications. The architecture supports asynchronous and synchronous communications using standard XML and JSON Representational state transfer (REST) message formats. Integration Server is multi-threaded, which allows for concurrent execution of multiple client requests. For secure client-to-server communication, Integration Server supports SSL. Third-party applications can use Integration Server to flexibly incorporate Perceptive Content functionality directly into products and enterprise systems by providing the required classes, functions, methods, subroutines, or other routines. Integration Server interacts seamlessly with Perceptive Content Server to increase efficiency with the use of your existing development tools.

Server Components

The Server Components include the Perceptive Content Server, Perceptive Content database, and the OSM.

Perceptive Content Server

The Perceptive Content Server can run on a variety of server-class machines ranging in configuration from economical to robust. You can run Perceptive Content in an enterprise environment on Microsoft Windows and Linux. The Perceptive Content Server requires standard TCP/IP networking to the clients, middle-tier products like Integration Server, the OSM, and the database.

Database

Perceptive Content customers can choose from Oracle or Microsoft DBMSs. Customers using the Perceptive Content Database Connector are free to house the database on the Perceptive Content Server or on a separate dedicated, relational database in order to centralize administrative functions.

The Perceptive Content Database Connector uses platform-specific drivers from industry-leader DataDirect to establish an ODBC connection between the Perceptive Content Server and the Perceptive Content database running in the RDBMS of your choice. This unique, “clientless” design takes advantage of specific queries and database features of the chosen platform and outperforms applications that use the native database (proprietary) API for database access.

OSM

Only the Perceptive Content Server accesses the OSM. The existing connection between Perceptive Content Server and the OSM is a standard file connection. When the server and the OSM are on the same machine, you have standard file connection on that machine. When the server and OSM are on different machines, you can access the OSM through a VPN that you set up to establish and maintain an encryption system through your network. Only the Perceptive Content Server accesses the OSM. The Perceptive Content Server accesses these files using standard operating system open and close calls.

You can cache files that are read or written to the main OSM storage device on a faster device, such as a SAN (Storage Array Network), to improve performance and reduce bottlenecks. The object store needs to be directly accessible by Perceptive Content Server. The object store for scanned images and other content can grow quite large.

You can also use a third-party hierarchical storage management (HSM) system to store OSM data. An HSM system stores objects in tiers, based on the frequency of access. Since the main purpose of object storage is to store fixed content in a form that is easy to access, an HSM system works well with the OSM. Be sure to consult with the HSM software documentation for recommendations and requirements on backup strategies.

In addition, when a customer chooses CAS (Content-Addressable Storage) devices over other storage options, such as optical disk or tape, Perceptive Content uses the CAS API to store and retrieve content from the CAS device.

Network Information

Perceptive Content components use packet delivery on the TCP/IP network. A packet is the unit of data that a network routes between an origin and a destination on the network. For example, when Perceptive Content Server sends information to the Perceptive Content Client, the TCP layer divides the file into chunks of an efficient size for routing.

You can configure the actual size of the packets using your network software. For example, you can change the packet size to the maximum size for your network. The network fills up the packet with data before sending it. The bigger the packet, generally, the more efficient the data transfer.

The Design Pillars of the Physical Perspective

Pillar	Description
Flexibility	**OSM Deployment** You can deploy the OSM on a SAN or a third-party hierarchical storage management (HSM) system, or other storage options, such as optical disk or tape. **Supported Operating Systems** The Perceptive Content Server runs on a variety of operating systems, including Microsoft Windows and Linux. **Supported DBMSs** Perceptive Content allows customers to choose from Oracle and Microsoft relational DBMSs.
Functionality	**Database Encryption (database supported)** Perceptive Content supports database encryption, which is a feature of the relational DBMS. **Transport Encryption** Perceptive Content can provide encrypted, port-level access to authenticated users via 256-bit AES encryption, which encrypts the Perceptive Content TCP/IP client and server socket communication. **Network Communications** To provide greater efficiency over various network topologies, Perceptive Software developed an efficient communications subsystem called “Tri-state Intelligent Sockets,” or TSIS™, to provide fast processing between the Perceptive Content Client pool and the Perceptive Content Server. TSIS bursts packets efficiently between the Perceptive Content Client pool and the Perceptive Content Server, enabling the Perceptive Content Server to service client requests in volumes beyond what would ever be experienced in the real world. The Perceptive Content TSIS ensures that Perceptive Content Clients experience no delays, optimizes load handling, and provides users with sub-second content retrieval.
Usability	**Security** The physical perspective of Perceptive Content shows an important usability factor, which is security. Your physical network provides security through features and encryption techniques available by the network, database, and OSM storage solution, as well as within the product security features of Perceptive Content.

Pillar	Description
Scalability	**64-Bit Support** To provide improved performance and additional memory capacity, you can install Perceptive Content Server in a 64-bit Windows environment. **OSM Deployment** You can use a redundant array of independent disks (RAID) level-5 system or an object-specific storage system to store the OSM, a third-party hierarchical storage management (HSM) system, or EMC Centera to store OSM data. All of these options are scalable as your capacity needs grow. **Clustering** For load-balancing, you can cluster the database server or OSM storage device on their respective servers. The hardware on which the database or OSM runs provides the clustering functionality. **DBMS** The Perceptive Content Server works with Oracle and SQL Server relational DBMSs. Each of these options is scalable, as proven by their manufacturer.
Extensibility	**Distributed Design** The core functionality the Perceptive Content Server delivers to end users can be centralized or distributed. This design enables your company to scan, process content, index, and view documents in one location or across remote locations. Additionally, you can distribute Perceptive Content Client components across servers in your enterprise, according to what best meets your needs and your environment.

Pillar	Description
Dependability	**Network Standards** Because Perceptive Content is a pure TCP/IP-based product, it offers the highest level of raw network performance and the greatest breadth of back-end scalability. **Backup** Perceptive Software recommends that you perform a full backup of the Perceptive Content object store and database on a nightly basis, or more often depending on your transaction load. For more information about maintaining your system and DBMS, refer to *Perceptive Content Best Practices for Oracle* or *Perceptive Content Best Practices for Microsoft SQL Server*. **Active-Active server environment** With Perceptive Content 6.7, you can install and configure Perceptive Content Server as an active-active environment (also known as a high-availability environment). Perceptive Content Server is installed as an instance – with copies of files, databases and so on – on a physical server. In its simplest form, an active-active server environment has at least two Perceptive Content instances running at the same time, a primary instance and a secondary instance; however, each instance of the Perceptive Content Server runs independently of the other. If the primary environment fails, there is no interruption when the system immediately switches over to the secondary environment without the need for a server-application restart. **Failover Support** A failover server can be particularly beneficial to customers using Perceptive Content across their enterprise in support of critical business processes. With the execution of a failover license agreement, you can create a failover instance of the Perceptive Content Server. In addition, Perceptive Content supports the ability to configure both a redundant object store and a redundant metadata repository, to ensure 100% data availability in case of a media failure. You can mirror the OSM to provide failover support in the event the file system for the primary OSM fails, while you use third party data/server mirroring tools to mirror all server resource and document objects to a separate server instance in the event the primary Perceptive Content Server goes down. Customers can take advantage of relational DBMS replication capabilities to create a real-time, redundant metadata repository on another server instance. **Load Balancing and Failover** The Perceptive Content Server Agents and companion products are remotable, meaning that you can install them independently on other servers. This achieves efficient load balancing across multiple CPUs and protects processing bandwidth on the Perceptive Content Server. For example, the database can be loaded on one server, the user load on another server, and the Perceptive Content supporting agents on one or more additional servers. **High Availability** Perceptive Content uses a pure IP transport between client and server. An IP connection from a client application to its corresponding server application is a key requirement for any application to server to work successfully in a high-availability architecture. Because the high-availability tool transparently migrates the IP address of the application service to whichever node it is currently running on, Perceptive Content can be transparently accessed on any server in a cluster. Perceptive Content requires no special configuration options outside of most high-availability tools. The various Perceptive Content services that are part of the cluster are simply managed in that cluster using standard administrative tools.

Perceptive Content Architecture: a Security Perspective

The Perceptive Content security features are designed to be configurable based on your organization's specific business rules, which allows you to give authorized users appropriate access to documents and functions when and where they need them. Perceptive Content content transmission is secured by high-end encryption that is only delivered to users who have been granted appropriate access privileges.

The intuitive graphical interface of the Perceptive Content Client allows you to carefully establish document, user, and group security with simple clicks of the mouse. Perceptive Content Client features, such as built-in administrative roles, multiple group membership, specific Global privileges, management-level privileges, and business application security make targeted control easy to establish and maintain.

Content Security

Some document management, imaging, and workflow systems require users to map to the volumes that store scanned and other object data, which not only circumvents basic security guidelines, but creates a much more complicated user management environment as well. In contrast, because of the Perceptive Content TCP/IP-based client/server architecture, the only way to access an object store managed by Perceptive Content is by logging onto the Perceptive Content Server. Eliminating file mapping dramatically improves performance and enhances security by eliminating the ability of users to traverse the file system where objects are stored.

Perceptive Content provides comprehensive security for content, including:

1. **Batch Users**. This feature secures batches (content that is captured and processed in a group) by specifying the users or groups who can access, create, quality assure, and index within existing batches.

- **Document Type Privileges**. These privileges provide security so that only authorized users and groups can access documents based on the associated metadata. After a user has the ability to access a document, you can control the actions he or she can perform by the type of document itself.
- **Drawer Privileges**. These privileges provide security so that you can determine which drawers users and groups can access. Within a drawer, you grant Documents, Explorer/Project Viewer, Viewer, Document Management, and Batch privileges.
- **Global Privilege**. These privileges secure Search, Projects, Capture, Batch (General), Viewer (Unlinked Documents), and Manage privileges.
- **Process Privileges**. These privileges provide specific roles to increase security surrounding the management and usage of workflow processes.
- **Project Privileges**. These privileges give users and groups the ability to change or modify existing project types.
- **Group Membership**. This feature enables Department Managers to add users to a group through the Users pane.

Digital Signatures

Perceptive Content provides the capabilities to digitally sign electronic documents using Public Key Infrastructure (PKI) technology that is in compliance with the NIST Digital Signature Standards (DSS) published as FIPS 186-3. When using digital signatures on documents, you provide an increased assurance that the person approving the information is indeed the appropriate person. Digital signatures authenticate the signer of the document and verify the data integrity (contents) of the signed document. All digitally signed and legally binding information is maintained within Perceptive Content and available for viewing at a moment's notice.

Administering Security in Perceptive Content

Perceptive Content includes two built-in roles — Perceptive Manager and Department Manager. Each role incorporates a set of pre-defined administrative privileges, allowing you to quickly and easily create administrative Perceptive Content users.

The Perceptive Manager Role

The Perceptive Manager role administers two general aspects of the system that are inaccessible to other users. The first is the ability to manage the entire pool of global users. Perceptive Managers create all users in the system, regardless of department. Perceptive Managers can modify the profile information for all users, assign users privileges in the Global category, and remove users from the system. A Perceptive Manager cannot assign his or her own privileges, nor can a Perceptive Manager remove other Perceptive Managers from the system.

A Perceptive Manager also controls the department management for the system. Perceptive Managers are the only users that create the departments in the system, modify the department name and displayed information, remove departments from the system, move content to another department, create and modify department names, promote other users to the Department Manager role, and demote other users from being a Department Manager.

The Department Manager Role

Users assigned the Department Manager role have the ability to assign non-Global privileges to other users, to manage users and groups in their department, and to perform virtually every privilege in their own department(except to create and demote other Department Managers). Perceptive Content allows you to optionally assign the role of Department Manager to as many users as required, as a useful way to further distribute administrative duties. Only users assigned the Perceptive Manager role can deactivate other users.

Establishing Users and Groups

Perceptive Content supports the ability to create unlimited users and groups that map to any access needs you identify. For example, you can further distribute administrative responsibilities by creating groups designed to establish power users who can collectively create other groups, add users, and assign privileges. In addition, Perceptive Content supports streamlined administration with its ability to assign users to multiple groups, which allows administrators to assign the bulk of privileges at the group level, then assign user-specific privileges to support strategic security decisions (such as, to allow a specific user in a group the ability to delete documents).

Before users can perform any action in Perceptive Content, the Perceptive Manager or Department Manager must grant them privileges to do so. Privileges can be explicitly granted, denied, or left unstated.

Authentication and Client Validation

Perceptive Content requires user and password authentication against a directory service, such as LDAP or Microsoft Active Directory, a local machine, or the server platform operating system. In addition, Perceptive Content also supports authentication against a SQL table structure. Rather than storing user passwords, Perceptive Content checks to make sure the user name and password offered at sign-on identically match the user name and password stored in the external directory service.

In addition, you can configure Perceptive Content to support client validation. For example, Perceptive Content can authenticate a user based on his or her domain log on, bypassing the Perceptive Content

sign-on process if the user has authenticated to the domain with a valid user ID known by the Perceptive Content Server.

Encryption

Perceptive Content delivers AES (Advanced Encryption Standard) encryption using the AES-256 block cipher. This encryption method uses a 256-bit encryption key to encrypt plain text into cipher text. AES is the first publicly accessible and open cipher to gain the approval of the National Security Administration (NSA) for the encryption of top secret information. Perceptive Content's AES encryption uses FIPS 140-2 Certified Cryptographic Provider modules from Microsoft and Certicom. You can use AES for user authentication requests for Perceptive Content.

Whether capturing a document or streaming content to the client for retrieval, Perceptive Content can provide encrypted, port-level access to authenticated users through 256-bit AES encryption, which encrypts the TCP/IP message stream between the Perceptive Content Server and external agents, Perceptive Content Clients, Interact for Office Clients, and Interact for eCopy Clients. AES Encryption is optional and is not enabled by default. To enable encryption, you configure the inow.ini file on your Perceptive Content Server.

HTTP(S) (SSL) is a fully supported transport encryption method between Integration Server and REST clients. Perceptive Content offers 128-bit SSL encryption between Integration Server and any application using the Integration Server API. In addition, 128-bit SSL encryption is available for Interact clients

Auditing

You can configure multiple audit templates to log more than 500 different Perceptive Content actions, such as who viewed, deleted or re-indexed an object, where the user connected from, or when the user performed the action and the user's session information.

You can create an unlimited number of audit templates. You can designate specific users and groups to each audit template. While an audit template is turned on, Perceptive Content logs the audit of user activity to either the Perceptive Content database or an XML file. If you store your audit activity in the Perceptive Content database, you can run activity reports using the Business Insight suite. If you export user activity to an XML file, Perceptive Content exports a log file for each 24-hour period during which a user member logged onto the system. After Perceptive Content generates an XML file, you can review the file auditing data using the tool of your choice.