

kafka a very short introduction very short introd Academic PDF

kafka a very short introduction very short introd Kafka is a powerful and widely adopted open-source platform designed for building real-time data pipelines and streaming applications. As the digital landscape evolves rapidly, organizations are increasingly relying on Kafka to manage large volumes of data efficiently, reliably, and in real-time. This article provides a comprehensive yet succinct overview of Kafka, its core concepts, architecture, use cases, and benefits, helping both beginners and seasoned professionals grasp its significance in modern data ecosystems.

What Is Kafka?

Kafka, originally developed by LinkedIn and later donated to the Apache Software Foundation, is a distributed event streaming platform. It allows applications to publish, process, and subscribe to streams of records in a fault-tolerant manner. Unlike traditional messaging systems, Kafka is designed to handle high-throughput, scalable, and durable data streams, making it ideal for real-time analytics, log aggregation, and event sourcing.

The Core Concepts of Kafka

Understanding Kafka's fundamental components is crucial to appreciating its capabilities. Here are the key concepts:

1. Topics

Topics are logical channels in Kafka where data records are published. They serve as categories or feeds for different data streams. Producers send data to topics, and consumers subscribe to topics to receive data.

2. Producers

Producers are client applications that publish data to Kafka topics. They determine which topic to send data to and handle the serialization of records.

3. Consumers

Consumers subscribe to topics to read data. They can read data from one or multiple topics and process it as needed. Kafka supports consumer groups, enabling load balancing and parallel

processing.

4. Partitions

Each Kafka topic is divided into partitions, which are ordered, immutable sequences of records. Partitions allow Kafka to distribute data across multiple brokers, facilitating scalability and fault tolerance.

5. Brokers

Kafka runs on a cluster of servers called brokers. Each broker manages a subset of partitions and ensures data is stored and replicated efficiently.

6. Replication

To ensure durability and fault tolerance, Kafka replicates partitions across multiple brokers. If one broker fails, data remains accessible from replicas.

Kafka Architecture Overview

Kafka's architecture is designed for high performance and resilience. Here's an outline of its main components and how they interact:

1. Kafka Cluster

A Kafka cluster comprises multiple brokers working together. This distributed setup allows Kafka to handle large-scale data streams seamlessly.

2. Producers and Consumers

Producers send data to Kafka, while consumers retrieve data for processing or storage. Both are decoupled, enabling flexible data pipelines.

3. Topics and Partitions

Data flows through topics, which are partitioned to facilitate parallel processing. Producers decide how to distribute data across partitions, often using keys for ordering.

4. Zookeeper

Zookeeper manages Kafka's metadata, including broker coordination, topic configuration, and

partition leadership election. (Note: Recent Kafka versions are moving towards removing Zookeeper dependency with KRaft mode.)

Use Cases of Kafka

Kafka's versatility makes it suitable for a wide range of applications across industries:

- **Real-Time Analytics:** Processing streaming data for dashboards, alerts, and business insights.
- **Log Aggregation:** Collecting logs from multiple servers and applications for centralized analysis.
- **Event Sourcing:** Capturing all changes to system state as a sequence of events.
- **Stream Processing:** Enabling complex transformations and computations on live data streams using tools like Kafka Streams or ksqlDB.
- **Data Integration:** Connecting different data systems, databases, and applications through real-time data pipelines.

Advantages of Using Kafka

Kafka offers numerous benefits that have contributed to its widespread adoption:

- **High Throughput:** Capable of handling millions of messages per second with minimal latency.
- **Scalability:** Easily scaled horizontally by adding more brokers and partitions.
- **Fault Tolerance:** Data replication ensures durability even if some brokers fail.
- **Durability:** Data is stored persistently on disk, reducing the risk of data loss.
- **Flexibility:** Supports multiple messaging patterns, including publish/subscribe and queue-based messaging.
- **Integration:** Compatible with various data processing frameworks and programming languages.

Getting Started with Kafka

For those interested in exploring Kafka, here are the basic steps:

1. Installation

Kafka can be run locally for development purposes or deployed on cloud infrastructure for production. The official website provides detailed installation guides.

2. Creating Topics

Use Kafka's command-line tools or APIs to create topics suited to your data streams.

3. Producing Data

Develop producer applications in languages like Java, Python, or Go to publish data to Kafka topics.

4. Consuming Data

Build consumer applications to subscribe to topics and process incoming data.

5. Monitoring and Management

Use Kafka's management tools and metrics to monitor cluster health, throughput, and lag.

Future of Kafka

Kafka continues to evolve, with ongoing developments focusing on simplifying management, improving scalability, and reducing dependencies. Projects like KRaft mode aim to replace Zookeeper, making Kafka more lightweight and easier to operate. Additionally, Kafka's ecosystem includes tools like Kafka Connect for data integration, Kafka Streams for stream processing, and ksqlDB for SQL-based stream querying, broadening its applicability.

Conclusion

Kafka stands out as a robust, scalable, and versatile platform for real-time data streaming. Its architecture supports high throughput, fault tolerance, and flexible deployment options, making it an essential tool for modern data-driven organizations. Whether used for real-time analytics, log aggregation, or building event-driven systems, Kafka empowers businesses to harness their data streams effectively. As data continues to grow exponentially, Kafka's role in managing and processing high-velocity data streams is set to become even more critical in the digital age.

Kafka: A Brief Yet Profound Introduction to the Master of Modern Data Streaming

In the rapidly evolving landscape of data processing and real-time analytics, few technologies have left as indelible a mark as Apache Kafka. Conceived at LinkedIn in 2010 and subsequently donated to the Apache Software Foundation, Kafka has become the backbone of many large-scale, distributed, and fault-tolerant data pipelines. Its core promise centers around enabling organizations to process, store, and analyze streaming data with unparalleled speed and reliability. This article aims to delve deeply into Kafka's architecture, operational principles, and its significance in contemporary data ecosystems, offering a comprehensive review suitable for technologists, researchers, and enterprise stakeholders.

Understanding Kafka: A Concise Overview

Kafka is fundamentally a distributed event streaming platform designed to handle high-throughput, real-time data feeds. Unlike traditional message brokers, Kafka's architecture emphasizes scalability, durability, and low latency, making it suitable for applications ranging from log aggregation and event sourcing to complex data pipelines and real-time analytics.

At its core, Kafka functions as a distributed commit log, where data is appended sequentially by producers and read by consumers in an organized, fault-tolerant manner. Its design allows it to process millions of messages per second across clusters comprising thousands of nodes, all while maintaining message ordering and durability.

Architectural Foundations of Kafka

A thorough understanding of Kafka's architecture is essential to appreciate its robustness and flexibility. The system is built around several key components:

Topics and Partitions

- Topics: Logical channels used to categorize and organize data streams.
- Partitions: Subdivisions within topics that enable parallelism and scalability. Each partition is an immutable sequence of messages, ordered and stored separately.

Producers and Consumers

- Producers: Clients that publish data to Kafka topics.
- Consumers: Clients that subscribe to topics and process incoming data.

Broker Cluster

- Kafka operates as a cluster of brokers, each managing a subset of topic partitions.
- Brokers coordinate to distribute load and ensure fault tolerance.

Replication and Fault Tolerance

- Each partition can be replicated across multiple brokers.
- Replicas ensure data durability and high availability, even in the event of broker failures.

ZooKeeper Integration

- Kafka relies on Apache ZooKeeper for cluster coordination, including broker registration, leader election, and configuration management.
- Note: Recent developments aim to reduce Kafka's dependency on ZooKeeper, moving toward a KRaft mode.

Operational Mechanics and Data Flow

Understanding how Kafka handles data is crucial to leveraging its capabilities:

- Data Production: Producers send messages to specific partitions within topics, often using partitioning strategies based on message keys for ordering guarantees.
- Data Storage: Messages are stored sequentially in partition logs, with each message assigned an offset—a unique identifier within the partition.
- Data Consumption: Consumers subscribe to topics and fetch messages based on offsets, allowing for scalable and independent processing.
- Retention Policies: Kafka retains data based on time or size limits, supporting use cases like replay or historical analysis.

Kafka's Ecosystem and Tooling

Beyond core functionality, Kafka's ecosystem includes a rich set of tools and components:

- Kafka Connect: Facilitates integration with external systems such as databases, data warehouses, and file systems.
- Kafka Streams: A Java library enabling real-time stream processing within Kafka.
- KSQL: A SQL-like interface for stream querying and transformation.
- Confluent Platform: An enterprise-grade distribution offering additional tools, management, and security features.

Use Cases and Industry Adoption

Kafka's versatility has led to widespread adoption across diverse industries:

- Financial Services: Real-time fraud detection, transaction monitoring, and market data processing.

- E-Commerce: Personalization, order tracking, and inventory management.
- Telecommunications: Call detail records, network monitoring, and customer analytics.
- Healthcare: Patient data streaming, device telemetry, and compliance auditing.
- Technology Companies: Log aggregation, metrics collection, and event sourcing.

Key use cases include:

- Event Sourcing and CQRS Architectures
- Real-Time Analytics and Dashboards
- Data Integration and ETL Pipelines
- Monitoring and Alerting Systems
- Machine Learning Data Feeds

Advantages and Challenges of Kafka

While Kafka offers compelling benefits, it also presents challenges that organizations must consider:

Advantages

- High throughput and low latency
- Horizontal scalability
- Fault tolerance and durability
- Flexibility in data processing paradigms
- Rich ecosystem and community support

Challenges

- Operational complexity, especially at scale
- Dependency on ZooKeeper (though evolving)
- Learning curve for designing optimal architectures
- Managing data retention and storage costs

Future Directions and Trends

Kafka continues to evolve, with recent initiatives focusing on:

- KRaft Mode: Eliminating ZooKeeper dependency for simpler deployment.

- Tiered Storage: Integrating with cloud storage solutions for cost-effective retention.
- Enhanced Security: Improving encryption, authentication, and authorization mechanisms.
- Stream Processing Enhancements: Expanding capabilities with more powerful, user-friendly processing tools.

As organizations increasingly rely on real-time data, Kafka's role is poised to grow, integrating more deeply with cloud-native architectures, serverless computing, and AI-driven analytics.

Conclusion

Kafka exemplifies a paradigm shift in data processing, transforming how organizations handle streaming data. Its architecture, built on distributed logs, high scalability, and fault tolerance, addresses the demands of modern data-driven applications. While operational challenges exist, the ecosystem's richness and continuous development ensure Kafka remains at the forefront of real-time data streaming solutions. Understanding its core principles, architecture, and strategic applications provides invaluable insights for anyone aiming to harness the power of streaming data in the digital age.

QuestionAnswer What is Kafka in simple terms? Kafka is an open-source distributed streaming platform used for building real-time data pipelines and streaming applications. Why is Kafka considered a key technology in data architecture? Because it allows high-throughput, fault-tolerant, and scalable handling of real-time data streams across systems. What are the main components of Kafka? Kafka's main components include producers, brokers, topics, partitions, and consumers. How does Kafka ensure data durability? Kafka stores data on disk and replicates it across brokers to prevent data loss and ensure durability. What is a Kafka topic? A topic is a category or feed name to which messages are published and from which they are consumed. Can Kafka handle real-time data processing? Yes, Kafka is designed for real-time data ingestion, processing, and analytics at scale. Is Kafka suitable for large-scale applications? Absolutely, Kafka is built for high scalability and can handle millions of messages per second. What is the difference between Kafka and traditional messaging systems? Kafka is distributed, scalable, and designed for high-throughput streaming, unlike traditional messaging queues which may be less scalable. How does Kafka support fault tolerance? Kafka replicates data across multiple brokers, ensuring data availability even if some brokers fail.

Related keywords: Kafka, Apache Kafka, message broker, event streaming, distributed system, pub/sub, real-time data, data pipeline, Kafka topics, Kafka architecture

Oracle Fnd Documents Short Text

oracle fnd documents short text is a commonly searched term by Oracle E-Business Suite (EBS) users and administrators seeking concise information about FND (Foundation) documents within the Oracle environment. Understanding how to efficiently access, interpret, and manage these short texts is crucial for maintaining seamless application functionality, troubleshooting issues, and ensuring effective communication across modules. This article provides an in-depth exploration of Oracle FND documents short text, covering its purpose, structure, usage, and best practices for managing these essential components.

Understanding Oracle FND Documents Short Text

What Are FND Documents?

In Oracle E-Business Suite, FND (Foundation) documents are standardized pieces of data used across various modules to support configuration, messaging, and processing. These documents often contain essential information such as error messages, labels, descriptions, and short texts that facilitate user interaction and system operations.

The Role of Short Texts in FND Documents

The short text within FND documents serves as a concise, easily understandable snippet of information that summarizes or labels a particular message or data element. This short text is typically used in user interfaces, notifications, and logs to quickly convey the essence of the message without overwhelming users with lengthy descriptions.

Key purposes of FND document short texts include:

- Providing quick, readable labels for form fields or options.
- Summarizing error or informational messages.
- Facilitating localization by enabling easy translation of concise texts.
- Enhancing user experience by reducing clutter and improving clarity.

Structure and Components of FND Documents Short Texts

Basic Structure of FND Documents

FND documents are stored within Oracle EBS's database tables, primarily in the `FND\_DOCUMENTS` and `FND\_DOCUMENTS\_TL` tables. These tables contain core information and translations, respectively.

Key tables involved:

- `FND\_DOCUMENTS`: Stores document identifiers, types, and core attributes.
- `FND\_DOCUMENTS\_TL`: Stores translated or localized text for different languages.

Components of Short Texts

A typical FND document short text consists of:

- Document Name/ID: Unique identifier for the document.
- Language Code: Specifies the language for the short text.
- Short Text Content: The concise message or label.
- Description: Additional details, often optional.
- Effective Dates: Validity period of the short text.

Example:

Column Name	Description
-----	-----
DOCUMENT_NAME	Unique identifier like 'ERR_INV001'
LANGUAGE_CODE	'US' for English (United States)
SHORT_TEXT	'Invalid input'
DESCRIPTION	'Error message for invalid user input'
START_DATE	Date from which the text is valid
END_DATE	Date until the text is valid

Usage of FND Document Short Texts in Oracle EBS

Application in User Interface

Short texts are used extensively in forms, menus, and notifications to:

- Display labels for fields.
- Show error messages or warnings.

- Present status updates or informational messages.

Example: When a user enters invalid data, the system may display a short text like "Input error" to communicate the issue promptly.

Application in Messaging and Logging

Short texts facilitate quick identification of messages in system logs, aiding in troubleshooting and auditing.

Localization and Multilingual Support

Because short texts are stored with language codes, they support localization efforts, enabling the interface and messages to adapt to users' language preferences.

Managing FND Documents Short Texts

Creating and Updating Short Texts

To create or modify short texts, administrators typically:

- Use Oracle Applications Developer or similar tools.
- Access the `FND\_DOCUMENTS\_TL` table directly via SQL for advanced management.
- Use the Oracle EBS System Administrator responsibilities to manage translations and texts.

Steps to create a new short text:

- Define a unique document name or ID.
- Specify the language code.
- Enter the concise message in the `SHORT\_TEXT` field.
- Set the effective date range.
- Save the record to make it available system-wide.

Best Practices for Managing Short Texts

- Consistency: Use standardized language and terminology across all texts.
- Clarity: Keep messages brief but informative.
- Localization: Ensure translations are accurate and contextually appropriate.

- Version Control: Track changes and maintain historical records for audit purposes.
- Validation: Regularly review short texts for relevance and correctness.

Common Tools and Techniques

- Application Developer: For creating and editing documents within the Oracle EBS interface.
- SQL Scripts: For bulk updates or extraction of short texts.
- FNDLOAD Utility: A command-line tool for exporting/importing document definitions, including short texts.
- Workflow and Notifications: Incorporate short texts into workflows for automated messaging.

Best Practices and Troubleshooting

Ensuring Data Integrity

- Regularly validate the consistency of short texts across different languages.
- Maintain proper date ranges to avoid displaying outdated messages.

Handling Missing or Incorrect Short Texts

- Verify if the document exists in `FND\_DOCUMENTS\_TL`.
- Check effective dates and language codes.
- Use SQL queries to identify missing translations, e.g.:

```
```sql
```

```
SELECT FROM FND_DOCUMENTS_TL
```

```
WHERE DOCUMENT_NAME = 'ERR_INV001'
```

```
AND LANGUAGE_CODE = 'US';
```

```
```
```

- Update or add missing entries as needed.

Performance Optimization

- Index the `FND\_DOCUMENTS\_TL` table on document name and language code.
- Clean up obsolete or unused texts periodically.

Conclusion

Oracle FND documents short text is a vital element in the Oracle E-Business Suite ecosystem, contributing significantly to user interface clarity, system messaging, and localization efforts. Proper management of these short texts ensures a more efficient and user-friendly environment, reduces errors, and enhances overall system performance. Whether you're a system administrator, developer, or functional consultant, understanding the structure, usage, and management of FND document short texts is essential for maintaining a robust Oracle EBS deployment.

Key Takeaways:

- Short texts are concise messages stored within FND documents.
- They support localization, usability, and troubleshooting.
- Proper management involves creation, updates, validation, and periodic review.
- Tools like FNDLOAD and SQL scripts facilitate efficient handling.
- Best practices ensure consistency, clarity, and system integrity.

By mastering the concepts and techniques outlined in this article, Oracle EBS users and administrators can optimize the use and management of FND document short texts, leading to improved system performance and user satisfaction.

Oracle FND Documents Short Text: An In-Depth Review and Guide

Oracle E-Business Suite's Foundation (FND) module plays a pivotal role in managing core system functionalities, including security, user management, and configuration. Among its myriad features, the handling and management of FND documents short text is a critical aspect that often puzzles users and administrators alike. This comprehensive review aims to demystify what FND documents short text entails, its significance, management practices, common issues, and best practices to optimize its use within Oracle EBS.

Understanding Oracle FND Documents Short Text

What is FND Documents Short Text?

FND Documents Short Text is a concise, descriptive field used within Oracle E-Business Suite to provide brief, meaningful summaries or titles for various documents, records, or transactions stored within the FND schema. It acts as a quick reference or identifier, allowing users to easily recognize

and differentiate documents without delving into detailed content.

Typically, this short text:

- Serves as a summary or caption for a document, record, or transaction.
- Is stored in the FND_DOCUMENTS table, primarily in the SHORT_TEXT column.
- Is used across multiple modules, especially in areas like approvals, notifications, and document management.

This field is designed to enhance usability by providing a human-readable, easily scannable label, which is essential in environments handling thousands of documents and records.

The Role and Significance of Short Text in Oracle FND

Why is Short Text Important?

The short text attribute plays several vital roles within the Oracle EBS environment:

- Quick Identification: It allows users to quickly identify the purpose or nature of a document without opening detailed records.
- Improved User Experience: Simplifies navigation and search processes, especially in approval workflows or when retrieving documents.
- Reporting & Auditing: Facilitates reporting by providing meaningful labels that can be used in queries, logs, and audit trails.
- Consistency & Standardization: Ensures that documents are labeled uniformly, aiding in compliance and process standardization.

Typical Use Cases

- Document Management: Short text labels for uploaded or generated documents such as invoices, receipts, or correspondence.
- Workflow Approvals: Brief descriptions of approval requests or notifications.
- Error & Exception Tracking: Short descriptions of issues or alerts generated within the system.
- Custom Records: User-defined documents or records created via custom modules or extensions.

Managing FND Documents Short Text

Creating and Updating Short Text

Managing the short text effectively involves understanding where and how to set or modify this attribute:

- During Document Creation: When new documents are uploaded or generated via forms, APIs, or custom code, the short text is usually provided as part of the creation process.
- Via API Calls: Developers can set or update the short text using specific APIs like `FND_DOCUMENTS_PUB.Create_Document``.
- Manual Updates: Administrators or power users can update the short text directly through SQL scripts or through custom interfaces, with caution.

Best Practices for Managing Short Text

- Use Clear and Concise Descriptions: The short text should be meaningful yet brief, ideally not exceeding 50 characters.
- Maintain Consistency: Follow naming conventions or labeling standards across the organization.
- Update Regularly: When document content or context changes, update the short text to reflect the current status or description.
- Automate When Possible: Use APIs or scripts to populate short text fields during batch processing or integrations, reducing manual errors.

Common Methods to Handle Short Text

- Using SQL: Direct updates via SQL scripts (not recommended unless necessary and with proper backups).
- Using API: Preferred method for creating or updating documents programmatically.
- Using Forms: Oracle EBS forms or custom interfaces designed for document management.

Technical Aspects and Table Structures

FND_DOCUMENTS Table Overview

The core table where document-related data, including short text, is stored is FND_DOCUMENTS. Key columns include:

- DOCUMENT_ID: Unique identifier for each document.

- SHORT_TEXT: The brief description or title of the document.
- FILE_NAME: Name of the uploaded or stored file.
- DOCUMENT_TYPE: Type/category of the document.
- CREATION_DATE: When the document was created.
- LAST_UPDATE_DATE: Last modification timestamp.

Data Integrity and Constraints

- The SHORT_TEXT is typically constrained to a specific length (often 50 characters), ensuring standardization.
- Proper indexing on DOCUMENT_ID facilitates quick retrieval.
- Referential integrity is maintained via foreign keys to other tables like FND_DOCUMENTS_TL for multilingual support.

Challenges and Common Issues with FND Documents Short Text

1. Inconsistent or Vague Short Texts

- Caused by manual entry errors or lack of standards.
- Leads to difficulty in document identification and retrieval.

2. Length Limitations

- The character limit may restrict descriptive detail.
- Users may resort to abbreviations that can be ambiguous.

3. Multilingual Support Concerns

- When supporting multiple languages, short text must be maintained across translations.
- FND's multilingual tables (e.g., FND_DOCUMENTS_TL) are used, which require careful synchronization.

4. Performance Impact

- Excessively long or numerous short texts can affect query performance, especially in large environments.

5. Synchronization and Data Consistency

- Ensuring that short text updates are reflected across all relevant modules and reports.

Best Practices for Optimizing FND Documents Short Text

1. Standardization

- Develop organizational standards for naming conventions.
- Use templates or predefined phrases to ensure uniformity.

2. Validation

- Implement validation routines to prevent overly long or vague descriptions.
- Use check constraints or application-level validations.

3. Automation

- Leverage APIs for consistent and error-free updates.
- Automate batch updates for large document sets.

4. Multilingual Management

- Maintain translations in FND_DOCUMENTS_TL.
- Use consistent terminology across languages.

5. Regular Audits

- Periodically review documents' short texts for accuracy and relevance.
- Remove or archive outdated descriptions.

Integrating Short Text with Other Oracle EBS Components

Workflow and Approvals

- Short text is often displayed in approval notifications, dashboards, and worklists.
- Proper labeling improves decision-making efficiency.

Reporting and Analytics

- Short text fields are included in reports for summaries and summaries.
- They enable quick filtering and categorization.

Custom Development

- Developers should ensure that custom modules or integrations correctly handle the SHORT_TEXT attribute.
- When creating custom documents or records, always populate the short text for clarity.

Conclusion and Final Thoughts

The FND documents short text field, though seemingly simple, is a fundamental component of Oracle E-Business Suite's document management and communication ecosystem. Its proper management ensures that users can efficiently identify, search, and utilize documents, ultimately contributing to smoother workflows, better compliance, and more effective data management.

To leverage its full potential:

- Emphasize standardization and clarity.
- Use robust APIs and automation tools.
- Maintain multilingual consistency.
- Regularly audit and update descriptions.

By doing so, organizations can significantly enhance their operational efficiency and data clarity within Oracle EBS, avoiding common pitfalls and maximizing the value derived from their document management practices.

In essence, mastering the management of FND documents short text is crucial for administrators, developers, and end-users aiming for a streamlined, consistent, and effective Oracle E-Business Suite environment.

QuestionAnswer What are Oracle FND documents short text used for? Oracle FND documents short text are used to provide concise descriptions or summaries of various components, such as

profiles, messages, or setup instructions within the Oracle E-Business Suite, facilitating quick understanding and reference. How can I modify the short text for an Oracle FND document? You can modify the short text by accessing the relevant form or table within Oracle Applications, such as the System Profiles or Messages window, and updating the description or text fields accordingly. Where are Oracle FND documents short texts stored? Short texts for Oracle FND documents are stored within the database tables associated with the specific modules, such as FND_DEFINITION or FND_MESSAGES, allowing for easy retrieval and management. Can I customize the short text in Oracle FND documents for different environments? Yes, customization is possible by updating the short texts in the database or through the Application Developer tools, enabling environment-specific descriptions or instructions as needed. Are there best practices for managing Oracle FND document short texts? Best practices include maintaining consistency in terminology, documenting changes properly, and ensuring that short texts are clear, concise, and accurately reflect the purpose of the documents to improve user understanding.

Related keywords: oracle fnd documents, oracle fnd guide, oracle fnd user manual, oracle fnd documentation, oracle fnd help, oracle fnd reference, oracle fnd tutorials, oracle fnd setup, oracle fnd administration, oracle fnd support

Read the full article online: [ORACLE FND DOCUMENTS SHORT TEXT](#)