

Data Analytics: Data Mesh Architecture and Its Implications for Data Management

Kishore Gade

Vice President, Lead Software Engineer at JP Morgan Chase

Abstract: Data mesh architecture is revolutionizing the field of data analytics, offering a fresh approach to how organizations manage and utilize data. Traditionally, centralized data warehouses and data lakes have dominated the landscape, but as the volume and complexity of data grow, these systems struggle to keep pace. Data mesh breaks away from the centralized model by promoting a decentralized and domain-oriented approach, where individual teams take ownership of their data. This shift enables organizations to treat data as a product, with each domain team responsible for managing and serving their data assets. The goal is to make data more accessible and usable across the organization while reducing bottlenecks that occur in centralized systems. Implementing data mesh architecture has profound implications for data management. It encourages better data governance, scalability, and flexibility, while fostering innovation and collaboration. However, it also introduces challenges, such as ensuring consistency, standardization, and proper integration across domains. The success of a data mesh strategy depends on a strong cultural shift, requiring both technical infrastructure and organizational alignment to support autonomous teams. Ultimately, data mesh offers the promise of unlocking the full potential of data in large organizations by making it more agile, democratized, and efficient. By embracing this architecture, businesses can better respond to the increasing demands of modern data management and gain a competitive edge in the digital landscape.

Keywords: Data Mesh, Data Management, Data Governance, Decentralized Data Architecture, Domain-Driven Design

1. Introduction

The rapid growth of data over the past decade has presented businesses with both incredible opportunities and serious challenges. As organizations generate and collect massive amounts of data, the task of managing, processing, and utilizing that data has become increasingly complex. Traditional methods, like data warehouses and data lakes, once considered the gold standard for data management, are now struggling to keep pace. These centralized systems often create bottlenecks, slow down decision-making, and reduce the flexibility that businesses need to stay competitive.

This is where data mesh architecture enters the picture. It's a relatively new approach that moves away from the monolithic, one-size-fits-all data management models toward something more distributed and adaptable. Instead of relying on a single, centralized repository for all data, data mesh advocates for a decentralized model, where data is owned and managed by individual teams, or "domains," within an organization. This approach aligns with the structure of modern businesses, which are often organized around different departments or product lines, each with its own data needs and expertise.

By treating data as a product and assigning responsibility for its management to the people closest to it, data mesh promotes greater autonomy, accountability, and scalability. Teams that own the data are better equipped to ensure its quality, make updates, and respond quickly to business demands. The shift to data mesh architecture also addresses some of the key pain points associated with traditional systems, such as the difficulties in scaling and the slow time-to-insight, which can stifle innovation and agility in decision-making.

In this article, we'll dive deeper into the concept of data mesh architecture, comparing it to traditional data management methods. We'll explore its potential to transform how

businesses handle data by fostering more effective governance, improving data ownership practices, and enabling greater scalability. Additionally, we'll look at how some companies have successfully implemented data mesh and discuss the cultural and operational changes necessary for making this shift work in practice. By the end, you'll have a clearer understanding of how data mesh can not only solve many of the problems plaguing traditional data management systems but also provide a future-proof strategy for leveraging data in an increasingly complex landscape.

2. The Evolution of Data Management: From Centralized to Decentralized Systems

2.1 Traditional Data Management Architectures

In the world of traditional data management, centralized architectures like data warehouses and data lakes have long been the foundation. These systems collect information from various sources, process it through transformations, and then store it in a single, centralized location. From there, different departments or business units within the organization can access and analyze the data to gain insights.

While this model works well for structured data and relatively simple analysis tasks, it has its share of limitations, especially as data volumes increase and business needs become more complex. Here are a few of the main challenges that organizations face when relying on centralized data architectures:

- **Scalability Issues:** As the amount of data grows exponentially, scaling up centralized systems to handle the increasing volume becomes a costly and complex endeavor. These systems often require more resources to store, process, and manage the data, leading to higher infrastructure costs.
- **Bottlenecks:** Centralized systems are typically managed by a single team, often referred to as the data engineering

or IT department. These teams can quickly become overwhelmed with requests from various departments, causing delays in the availability of the data. Business units may experience slowdowns in getting the information they need, which can impact decision-making and operational efficiency.

- **Lack of Domain-Specific Expertise:** A centralized data team, while technically proficient, may not always have the detailed knowledge required to fully understand the context or specific nuances of data generated by different business units. This lack of domain expertise can lead to data being managed in ways that don't fully align with the unique needs of individual departments, ultimately limiting the data's usefulness.

Because of these challenges, traditional centralized data architectures can often act as roadblocks to innovation, flexibility, and scalability. As data volumes continue to surge and the demand for real-time insights increases, organizations are beginning to rethink how they structure their data management systems. Enter: the concept of the data mesh.

2.2 The Emergence of Data Mesh

In response to the limitations of traditional centralized architectures, data mesh has emerged as a modern approach to data management. Unlike traditional models, data mesh is a decentralized architecture that distributes the responsibility for data management across the entire organization. It reimagines data as a product, giving individual teams ownership over their data, and ensures that each team is accountable for the quality, usability, and accessibility of the data they produce.

Here are the core principles of the data mesh approach:

- **Domain-Oriented Data Ownership:** One of the key tenets of data mesh is that the teams that generate the data are also responsible for managing it. This domain-oriented model assigns ownership to the teams that are closest to the data, making them responsible for ensuring its accuracy and relevance. For example, the marketing team would manage its own customer data, while the finance team would manage its financial data. This shift in ownership ensures that data is handled by those who understand its context and can make better decisions about how it's used.
- **Data as a Product:** In a data mesh architecture, data is treated like a product, which means it must meet certain standards of quality, usability, and availability. Each team is responsible for ensuring that their data is reliable and accessible to others in the organization who may need it. This approach encourages teams to view their data as an asset that can be shared and used across the organization, similar to how a well-designed product would be. By treating data as a product, organizations can improve data quality, ensure consistency, and make data more valuable.
- **Self-Serve Data Infrastructure:** To enable decentralized data management, teams need access to the right tools and platforms. A self-serve data infrastructure empowers individual teams to manage their own data without relying on a central IT or data engineering team for every request. With standardized tools and technologies, each team can store, process, and analyze their own data in a way that aligns with their specific needs. This self-service model

reduces bottlenecks, speeds up data access, and increases agility within the organization.

- **Federated Data Governance:** While data mesh promotes decentralization, it doesn't mean that governance is ignored. In fact, data mesh adopts a federated governance model where oversight and compliance responsibilities are distributed across different teams, but coordinated centrally. This ensures that data security, privacy, and compliance requirements are met without stifling innovation or autonomy. Teams have the freedom to manage their data, but they must also adhere to organization-wide governance policies and standards.

2.3 The Benefits and Implications of Data Mesh

By adopting a data mesh architecture, organizations can overcome many of the challenges posed by traditional centralized data systems. Here's a look at some of the key benefits and implications:

- **Increased Agility:** With teams having direct control over their own data, they can quickly access, analyze, and act on it without waiting for a centralized team to process requests. This can lead to faster decision-making and a more agile organization.
- **Scalability:** Since data is managed at the domain level, scaling becomes more efficient. Each team can scale their data management processes independently based on their needs, which reduces the complexity of scaling a massive, centralized system.
- **Improved Data Quality:** When teams are responsible for their own data and treat it as a product, they are more likely to ensure that it is accurate, complete, and well-maintained. This focus on quality can lead to better insights and more reliable decision-making across the organization.
- **Faster Time to Insights:** By decentralizing data management and reducing bottlenecks, data mesh allows teams to access the data they need more quickly. This faster access to data leads to quicker insights and improved business outcomes.
- **Empowerment of Domain Experts:** With data mesh, the teams that generate the data are also the ones using it for analysis. This empowers domain experts who have a deep understanding of the data to extract more meaningful insights, which can drive innovation and strategic decision-making.

However, it's important to recognize that implementing a data mesh architecture also comes with its own set of challenges. For instance, decentralized teams may need training and support to effectively manage their own data. Additionally, coordination across different teams for governance and compliance can be complex and require strong leadership.

3. Core Principles of Data Mesh Architecture

Data Mesh architecture represents a shift in how organizations think about and handle data. Instead of relying on a centralized approach, it decentralizes data management, ensuring that the people closest to the data are responsible for it. In this section, we'll break down the four core principles of data mesh architecture: domain-oriented decentralization, treating data as a product, self-serve data infrastructure, and

federated governance. These principles shape how data is managed, shared, and governed in modern organizations.

3.1 Domain-Oriented Decentralization

At the heart of data mesh is the idea of decentralizing data ownership. Rather than having a single team or department responsible for all of the organization's data, each business domain (or department) takes charge of its own data. This principle, known as domain-oriented decentralization, ensures that the people who understand the data best are the ones managing it.

For instance, in a retail company, the marketing department would manage customer data, while the operations team would handle inventory and supply chain data. This decentralization allows teams to work with data they are intimately familiar with and ensure it's being used in ways that align with their specific business goals. The decentralized approach also allows for more agility and flexibility. Changes in a specific domain's data needs can be addressed quickly, without having to go through a central data team that might have other priorities.

Additionally, this model fosters accountability within teams. When each domain owns and manages its data, they are responsible for ensuring the data's accuracy, relevance, and compliance with organizational and legal standards. As data is no longer the sole responsibility of a central IT or data team, it becomes a living, evolving part of the organization, directly tied to the goals and expertise of the people working with it.

3.2 Data as a Product

Another key principle of data mesh is the notion of treating data like a product. When data is viewed as a product, the teams responsible for it must ensure it's of high quality, easy to use, and accessible to others within the organization. Just as a company wouldn't launch a product without ensuring it meets customer needs, teams should approach data with the same level of attention.

What does this mean in practice? It means that data teams must think beyond merely storing and processing data. They should focus on making it clean, accurate, well-documented, and available in a way that others in the company can easily find and understand it. For instance, the marketing team responsible for customer data must ensure that this data is up-to-date, free from errors, and accompanied by clear documentation that explains how to use it.

Furthermore, this principle promotes a "customer-first" mindset, where the data producers—the teams managing the data—are accountable to the data consumers—the people who need to use the data in other departments. This focus on usability encourages continuous improvement of data quality and accessibility, fostering a culture of collaboration across departments.

In essence, data as a product shifts the role of data teams from being mere data custodians to becoming service providers. They create, maintain, and refine data products that others can

rely on for decision-making, driving innovation, and improving business processes.

3.3 Self-Serve Data Infrastructure

To make domain-oriented decentralization work effectively, data mesh architecture depends on self-serve data infrastructure. This principle refers to the creation of standardized, user-friendly tools and platforms that allow individual teams to handle their own data without having to rely heavily on a central data engineering team.

Imagine a scenario where a marketing analyst needs to extract customer data and run a specific analysis. In a traditional setup, this request might go through a central IT or data team, which could take days or weeks to fulfill. In a data mesh environment, the marketing team has the tools to access the data directly, run their analyses, and make real-time decisions.

The self-serve aspect empowers teams to work more autonomously while reducing bottlenecks. Teams can build, manage, and access their data pipelines, perform analytics, and maintain their data products without needing extensive technical expertise. This democratization of data tools is essential to the success of the data mesh, as it ensures that all domains, regardless of technical ability, can contribute to and benefit from the data ecosystem.

In a practical sense, self-serve data infrastructure includes a range of tools, from data storage solutions and processing frameworks to data cataloging and discovery platforms. These tools must be intuitive and accessible to users from all domains, enabling them to focus on deriving insights from data rather than getting bogged down by technical complexities.

3.4 Federated Governance

While decentralization offers significant advantages, it also poses a challenge: how do you maintain consistency, quality, and compliance across the organization? This is where the concept of federated governance comes into play.

Federated governance in data mesh provides a framework that balances autonomy with oversight. Each domain has the freedom to manage its data, but they must do so within a set of organizational standards and guidelines. These standards ensure that data across the organization remains consistent, secure, and compliant with external regulations like GDPR or industry-specific data laws.

For example, while each department may manage its own data, the organization might enforce company-wide rules on data security, privacy, and interoperability. This ensures that even as data is decentralized, it can still be easily integrated, shared, and understood across different domains. The goal is to prevent data silos while maintaining flexibility and responsiveness in individual departments.

Federated governance is typically achieved through collaboration between domain experts and centralized data governance teams. These teams work together to define data

standards, security protocols, and access controls that apply organization-wide. At the same time, federated governance allows each domain to adapt these standards to their unique data needs.

By implementing federated governance, organizations strike a balance between empowering teams to innovate with their data and ensuring that data remains trustworthy, secure, and compliant across the entire business.

4. Implications of Data Mesh for Data Management

Data mesh is a game-changer for how organizations handle and manage their data, offering both opportunities and challenges. In this section, we will explore the key implications that data mesh architecture has for data management, including its effects on scalability, ownership, data access, governance, and the difficulties organizations may face when adopting this model.

4.1 Improved Scalability

One of the most valuable aspects of a data mesh architecture is its ability to significantly improve scalability. Traditional data systems often rely on centralized databases and pipelines that must expand as the organization grows, leading to bottlenecks and increased operational costs. With data mesh, however, each domain—such as marketing, sales, or product—handles its own data. This decentralized approach eliminates the need for a single, monolithic system to support the entire organization.

By breaking down data ownership and management into smaller, more manageable parts, data mesh allows organizations to scale more efficiently. Instead of investing in massive centralized infrastructure, teams can focus on scaling their own data needs. This not only reduces complexity but also cuts costs associated with constantly upgrading and maintaining a central system. Teams are free to innovate and grow their data capabilities without being bogged down by limitations of a central infrastructure, allowing organizations to scale both their data and operations smoothly.

4.2 Enhanced Data Ownership and Accountability

Another key advantage of the data mesh model is the way it shifts data ownership and accountability to the teams that generate and use the data. In traditional, centralized architectures, it is often unclear who is responsible for data quality, which can lead to inconsistencies, errors, and data silos. Data mesh solves this issue by making data ownership explicit: each domain is responsible for managing its own data.

This decentralization drives accountability. Teams that own their data are more likely to ensure its accuracy, consistency, and timeliness, as they are directly impacted by the quality of the data they manage. This approach fosters a sense of responsibility, motivating teams to improve their data practices. As a result, organizations can expect better data quality overall, as the people closest to the data are responsible for maintaining it.

4.3 Data Democratization

One of the most transformative elements of the data mesh is its support for data democratization. In traditional data systems, access to data is often controlled by a central IT or data team, making it difficult for individual teams to get the information they need quickly. This leads to bottlenecks, delayed decision-making, and a reliance on data specialists to handle even simple queries.

Data mesh flips this model on its head. By giving individual teams the tools and infrastructure to manage their own data, it enables more direct access to data across the organization. This "self-serve" approach allows teams to explore, analyze, and use data independently, without waiting on approvals or assistance from a central data department.

This increased accessibility fosters faster decision-making and encourages innovation. Teams can experiment with new ideas, identify trends, and make data-driven decisions on the fly, all without the barriers imposed by traditional, siloed data systems. In this way, data democratization helps create a more agile and responsive organization.

4.4 Challenges in Implementing Data Mesh

While the benefits of data mesh are compelling, it's important to acknowledge the challenges that come with implementing this architecture. Transitioning to a decentralized data management model requires a shift in both culture and operations, and these changes can be difficult to navigate.

- **Organizational Change:** Moving from a centralized to a decentralized system requires more than just technical changes; it demands a cultural shift as well. Teams must learn to collaborate in new ways, sharing data across domains and taking on new responsibilities for data management.
- **Training and Expertise:** The success of a data mesh model depends on the ability of teams to manage their own data effectively. This means that team members will need training in data management practices, tools, and technologies. Not every organization has the expertise or resources to upskill their employees in this way.
- **Cost and Complexity:** While data mesh can reduce long-term costs related to scaling infrastructure, there are upfront investments that need to be considered. Implementing a data mesh requires new tools, technologies, and governance frameworks, which can be expensive and complex to set up.
- **Coordination Between Teams:** With decentralization comes the challenge of ensuring that data remains consistent and coordinated across teams. Teams must work together to ensure data standards are met, which can be difficult without proper processes and tools in place.

These challenges highlight the need for careful planning, strong leadership, and a phased approach to implementation. It's essential that organizations understand the potential roadblocks and invest in the right resources to ensure a smooth transition.

4.5 Implications for Data Governance

Data governance is a crucial consideration in any data architecture, and this is especially true with data mesh. In traditional centralized systems, governance is often managed by a single team that oversees data policies, compliance, and quality across the organization. With a data mesh, however, governance is distributed, requiring a new approach to ensure compliance and consistency.

In a data mesh architecture, each domain is responsible for governing its own data, but this does not mean that governance is entirely fragmented. The concept of federated governance comes into play, where global standards and policies are established at an organizational level, while the implementation of those policies is handled by individual domains. This ensures that organizations can maintain control over critical aspects like security, privacy, and data quality, even as governance responsibilities are distributed.

Federated governance allows for flexibility without sacrificing control. Organizations can set overarching policies that ensure compliance with regulations (e.g., GDPR or HIPAA), while giving individual teams the freedom to manage their data in a way that suits their needs. This balance of autonomy and oversight is key to making data mesh work effectively at scale.

4.6 Encouraging Innovation and Agility

One of the less discussed but equally important implications of data mesh for data management is how it fosters innovation and agility. By giving teams ownership of their data and the tools to manage it, organizations are effectively empowering employees to be more creative and proactive in their approach to data.

Teams no longer need to wait for a central data department to approve projects or provide access to datasets. Instead, they can quickly experiment, test hypotheses, and launch new initiatives. This ability to work independently encourages innovation, as teams are more likely to take risks and explore new ideas when they have control over their data.

At the same time, this increased agility allows organizations to respond more quickly to changes in the market, customer preferences, or emerging trends. In a data mesh environment, teams can pivot more easily, as they are not held back by outdated or rigid data systems.

5. Real-World Applications of Data Mesh

The rise of data mesh architecture is not just a theoretical development in data management; it has practical applications that have transformed the way organizations handle their data. By decentralizing data ownership and focusing on domain-driven design, companies can ensure that the data is more accurate, accessible, and aligned with specific business needs. Let's explore two real-world case studies—an e-commerce company and a financial institution—that demonstrate the tangible benefits of adopting data mesh.

5.1 Case Study 1: E-Commerce Company

In today's digital marketplace, data is the lifeblood of e-commerce businesses. With the continuous influx of customer data, transaction details, and inventory information, companies face the growing challenge of managing and making sense of this data quickly and accurately. One large e-commerce company found that its traditional, centralized data infrastructure was struggling to keep up with its expanding operations. Data teams were often overwhelmed, and insights from data analysis were slow to surface, leading to missed opportunities and inefficiencies.

To combat this, the company implemented a data mesh architecture, which fundamentally changed the way data was managed. Instead of having one central team responsible for all the data, ownership was distributed among the different departments or domains—each in charge of its own data sets. For example, the marketing team managed customer engagement data, while the sales team focused on transaction data.

This decentralization of data ownership had multiple benefits:

- **Improved Data Quality:** Since each team was responsible for its own data, there was greater accountability. Teams were able to focus on their specific domain data and ensure its accuracy, relevance, and quality. Previously, when a single team managed all data, it was easy for important nuances to be lost or for data quality issues to be overlooked.
- **Faster Insights:** The time it took to generate insights was drastically reduced. Since each domain team could access and manage its own data, they could generate insights without waiting for the central data team to process and deliver it. This allowed for faster decision-making and the ability to respond more quickly to customer demands and market trends.
- **Cross-Functional Collaboration:** The company also noted improved collaboration between teams. Because data mesh made it easier for different departments to understand their data within the context of their own operations, teams were able to share insights more effectively across the organization.

This transformation enabled the e-commerce company to move away from the bottlenecks associated with a centralized data approach, becoming more agile in an increasingly competitive market. In the end, data mesh provided the structure needed to support the company's rapid growth while ensuring data remained high-quality and actionable.

5.2 Case Study 2: Financial Institution

Financial institutions, with their massive stores of transactional, customer, and market data, are often at the forefront of data management innovation. One large financial organization was facing significant challenges in managing the vast amounts of data generated by its various business units. The centralized data warehouse it relied on was becoming an obstacle to innovation. Data requests from different departments were piling up, and the centralized team couldn't meet the demand quickly enough. On top of that, the

warehouse was becoming increasingly expensive to scale as the volume of data grew.

Recognizing these limitations, the institution decided to shift to a data mesh architecture. Here's how the transition worked and the key outcomes:

- **Scalability:** With the decentralized approach, each business unit became responsible for its own data, enabling the institution to scale its infrastructure more efficiently. Instead of relying on a single, central data warehouse, each unit could manage its own data infrastructure, which was more flexible and cost-effective. As a result, the institution was able to handle the ever-increasing volume of data without incurring the ballooning costs associated with scaling a monolithic system.
- **Data Accuracy and Timeliness:** The financial institution saw a marked improvement in data accuracy. Each business unit had a deeper understanding of its data, and because they were responsible for managing it, the teams were able to ensure the data was clean, accurate, and up-to-date. This was particularly important in the finance sector, where the accuracy and timeliness of data can make or break critical business decisions.
- **Empowerment of Business Units:** One of the most significant shifts was in the empowerment of individual business units. Rather than waiting for a centralized team to provide data access, each unit had direct control over the data that was most relevant to its operations. For example, the retail banking team could analyze customer behavior and transactional trends on its own, while the investment team could dive deeper into market data without delays. This led to more informed and timely decisions, driving business growth and improving customer satisfaction.
- **Improved Data Governance:** Although data ownership was decentralized, the financial institution maintained a strong governance framework. By embedding governance practices within each domain, the company ensured that data quality standards were met, while still allowing for flexibility in how data was managed.

6. The Future of Data Mesh Architecture

Data mesh is still an evolving concept, but its potential to reshape the future of data management is undeniable. As more companies begin to adopt this architecture, we are likely to witness ongoing developments in both the underlying infrastructure and the tools designed to support it.

One of the key drivers behind the growing interest in data mesh is the increasing complexity of data ecosystems within organizations. Traditional data management strategies often struggle to scale as companies grow, making it difficult to extract value from vast amounts of data. Data mesh offers a promising solution by decentralizing data ownership, allowing teams to manage their own data as a product while adhering to company-wide standards.

As technology continues to advance, especially in fields like artificial intelligence (AI) and machine learning (ML), the need for scalable, decentralized data systems will only become more pressing. AI and ML models thrive on large

amounts of diverse data. In a traditional centralized data management system, bottlenecks can occur, slowing down access to the necessary data and hindering the ability to make real-time decisions. Data mesh, with its decentralized approach, provides a way to meet the growing demand for more flexible and scalable data solutions.

In the coming years, we can expect a rise in specialized tools and platforms designed specifically for implementing and managing data mesh architectures. This will likely include advancements in automation, data governance, and integration capabilities. The goal will be to make data mesh easier to adopt for organizations of all sizes, reducing the complexity that currently comes with its implementation. Companies that can efficiently manage their data ecosystems in this way will gain a significant competitive advantage, as they will be able to make faster, more informed decisions based on high-quality, accessible data.

Furthermore, data mesh's impact will extend beyond just technological innovations. It will also influence organizational culture and collaboration. As teams take greater ownership of their data, there will be a shift toward more cross-functional collaboration. Data will no longer be siloed in specific departments but rather treated as a shared asset that can drive business outcomes across the board.

One of the challenges of this shift, however, will be ensuring proper data governance. With multiple teams managing their own data domains, organizations will need to establish clear governance practices to maintain data quality, security, and compliance. Fortunately, tools and frameworks for managing decentralized data governance are already emerging, and we can expect more robust solutions to be developed as data mesh continues to gain traction.

7. Conclusion

Data mesh architecture is a game-changer in how organizations approach data management. By moving away from the traditional, centralized models and distributing data ownership across different business domains, it allows for a more adaptable and scalable approach. This decentralization helps ensure that data is managed by the teams who are most familiar with it, leading to better data quality and accountability. However, like any new framework, implementing a data mesh isn't without its hurdles. It requires shifts in mindset, changes in roles and responsibilities, and may even demand new tools and technologies.

As the volume and complexity of data continue to grow, the shortcomings of traditional architectures become harder to ignore. In contrast, a data mesh helps organizations stay agile and efficient, allowing them to harness the power of data without getting bogged down by bottlenecks or single points of failure. It's an approach that supports both scalability and innovation, enabling companies to evolve with their data needs.

Adopting a data mesh architecture isn't just about improving data management — it's about cultivating a culture where data becomes a core asset, empowering teams across the organization to make smarter, faster decisions. In the long

run, this approach can lead to better business outcomes, increased innovation, and a stronger competitive edge in an increasingly data-driven world.

References

- [1] Sahni, Y., Cao, J., Zhang, S., & Yang, L. (2017). Edge mesh: A new paradigm to enable distributed intelligence in internet of things. *IEEE access*, 5, 16441-16458.
- [2] Gharaibeh, A., Salahuddin, M. A., Hussini, S. J., Khreishah, A., Khalil, I., Guizani, M., & Al-Fuqaha, A. (2017). Smart cities: A survey on data management, security, and enabling technologies. *IEEE Communications Surveys & Tutorials*, 19(4), 2456-2501.
- [3] Siddiqua, A., Hashem, I. A. T., Yaqoob, I., Marjani, M., Shamshirband, S., Gani, A., & Nasaruddin, F. (2016). A survey of big data management: Taxonomy and state-of-the-art. *Journal of Network and Computer Applications*, 71, 151-166.
- [4] Cignoni, P., Montani, C., Rocchini, C., & Scopigno, R. (2003). External memory management and simplification of huge meshes. *IEEE Transactions on Visualization and Computer Graphics*, 9(4), 525-537.
- [5] Jun, S. W., Liu, M., Lee, S., Hicks, J., Ankorn, J., King, M., ... & Arvind. (2015). Bluedbm: An appliance for big data analytics. *ACM SIGARCH Computer Architecture News*, 43(3S), 1-13.
- [6] Hossain, E., & Leung, K. K. (Eds.). (2007). *Wireless mesh networks: architectures and protocols*. Springer Science & Business Media.
- [7] Abram, G., & Treinish, L. (1995). An extended data-flow architecture for data analysis and visualization. *ACM SIGGRAPH Computer Graphics*, 29(2), 17-21.
- [8] Lichtensteiger, B., Bjelajac, B., Müller, C., & Wietfeld, C. (2010, October). RF mesh systems for smart metering: System architecture and performance. In *2010 First IEEE International Conference on Smart Grid Communications* (pp. 379-384). IEEE.
- [9] Zheng, F., Abbasi, H., Docan, C., Lofstead, J., Liu, Q., Klasky, S., ... & Wolf, M. (2010, April). PreData—preparatory data analytics on peta-scale machines. In *2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)* (pp. 1-12). IEEE.
- [10] Baumann, P., Mazzetti, P., Ungar, J., Barbera, R., Barboni, D., Beccati, A., ... & Wagner, S. (2016). Big data analytics for earth sciences: the EarthServer approach. *International journal of digital earth*, 9(1), 3-29.
- [11] Kang, T. W., & Hong, C. H. (2015). A study on software architecture for effective BIM/GIS-based facility management data integration. *Automation in construction*, 54, 25-38.
- [12] Lamahamedi, H., Szymanski, B., Shentu, Z., & Deelman, E. (2002, October). Data replication strategies in grid environments. In *Fifth International Conference on Algorithms and Architectures for Parallel Processing, 2002. Proceedings.* (pp. 378-383). IEEE.
- [13] Waharte, S., Boutaba, R., Iraqi, Y., & Ishibashi, B. (2006). Routing protocols in wireless mesh networks: challenges and design considerations. *Multimedia tools and Applications*, 29, 285-303.
- [14] Anawar, M. R., Wang, S., Azam Zia, M., Jadoon, A. K., Akram, U., & Raza, S. (2018). Fog computing: An overview of big IoT data analytics. *Wireless Communications and Mobile Computing*, 2018(1), 7157192.
- [15] Cox, M., & Ellsworth, D. (1997, August). Managing big data for scientific visualization. In *ACM siggraph* (Vol. 97, No. 1, pp. 21-38). MRJ/NASA Ames Research Center.