

FalconStor[®]



BEST PRACTICES FOR OPTIMIZING IT

THE OPTIMIZED DATA SERVICES MODEL

OPTIMIZING IT:

THE OPTIMIZED DATA SERVICES

ABSTRACT

In an effort to achieve greater efficiencies for IT investments, many IT departments leverage the Information Technology Information Library (ITIL) standards model when implementing solutions. Although standardization provides substantial operational benefits, traditional methodologies for critical data services fall short in meeting the stringent financial and budgetary objectives required in a downturned market. A better system is needed, one that optimizes IT services in an integrated and comprehensive manner.

The Optimized Data Services (ODS) model embraces advancements in server and storage technologies such as virtualization, data deduplication, continuous data protection, WAN optimization, thin provisioning, and artificial intelligence to enable a paradigm shift in IT to a more cost-effective service-oriented approach to information management. Data services are the main functions performed by an IT department in support of the applications which run the business:

- Provisioning Storage
- Protecting Data
- Disaster Recovery
- Rapid Backup and Recovery
- Technology Refresh/Data Mobility

The optimized model focuses on helping organizations make prudent choices in technology based on the benefits to the organization and its applicability within the existing IT environment, with no deference towards vendor relationships or internal political biases. The goal is to expose any unknown or hidden costs associated with traditional approaches and the overall effect on budgets and operational efficiency.

The result of implementing an ODS model within IT causes a shift from manual complex processes to automated operations. Implementation of the server and storage infrastructure changes from a buy, build, and tune approach to a reference architecture modular building blocks tied together with an intelligent abstraction layer which provides policy based automated management and monitoring of the environment, while enabling complete mobility of information.

The end results can be remarkable:

- Dramatic increases in operational efficiency
- Backup times reduced by 70% or more
- Recovery times improved by 95% or more
- Reduced risk and possibility of data loss
- Optimized utilization of storage and network infrastructure
- Improvements in DR capabilities and costs
- Reduced overall operating expenses of 30-80%

This document provides an outline for conducting an internal audit of an existing IT infrastructure, including operations and costs, to assess where optimization could be applied to data services. It outlines the best practices for implementing an optimized IT model to simplify IT and make it more cost-effective and operationally efficient.

TABLE OF CONTENTS

Overview	3
Traditional vs. Optimized Architecture	3
1. Traditional Architecture	3
2. Optimized Data Services Model Architecture	5
Business Objectives	6
Strategic Planning	7
IT Infrastructure Assessment	8
Designing a Service Level Agreement (SLA) Framework	9
Cost/Benefit Analysis	10
Conclusion	11

OVERVIEW

Our dependency on real-time information requires continuous access to applications and non-disruptive protection of data. As a result, there is an increased need for new more flexible data management techniques for IT data services. Since IT plays a critical role in every business, having a more efficient IT infrastructure and operations than the competition can provide a real competitive advantage.

Organizations must leverage technical innovations to capitalize on their cost advantages, and align IT to the business so that future improvements in technology can provide an immediate advantage. The ODS model proposes a best-of-breed approach where existing IT investments are leveraged together with new innovation in technology to enable a more flexible, cost-effective service-oriented solution for data management.

The ODS model provides immediate improvements in overall IT costs, and provides more stringent recovery point objectives (RPO) and recovery time objectives (RTO). ODS reduces reliance on tape for backup and recovery, automates disaster recovery (DR), and introduces intelligent abstraction to simplify and automate operations and technology refresh.

TRADITIONAL VS. OPTIMIZED ARCHITECTURE

TRADITIONAL ARCHITECTURE

Traditional data services focus on the management of individual elements within the datacenter. A buy, build, and tune approach is used for systems and storage. A tape-centric methodology is used for backup and recovery, and array-based technologies are leveraged for data replication and disaster recovery. The traditional model may or may not provide tiered storage services, and if it does, it typically relies on firmware the storage arrays or host-based software to provide the data movement functions.

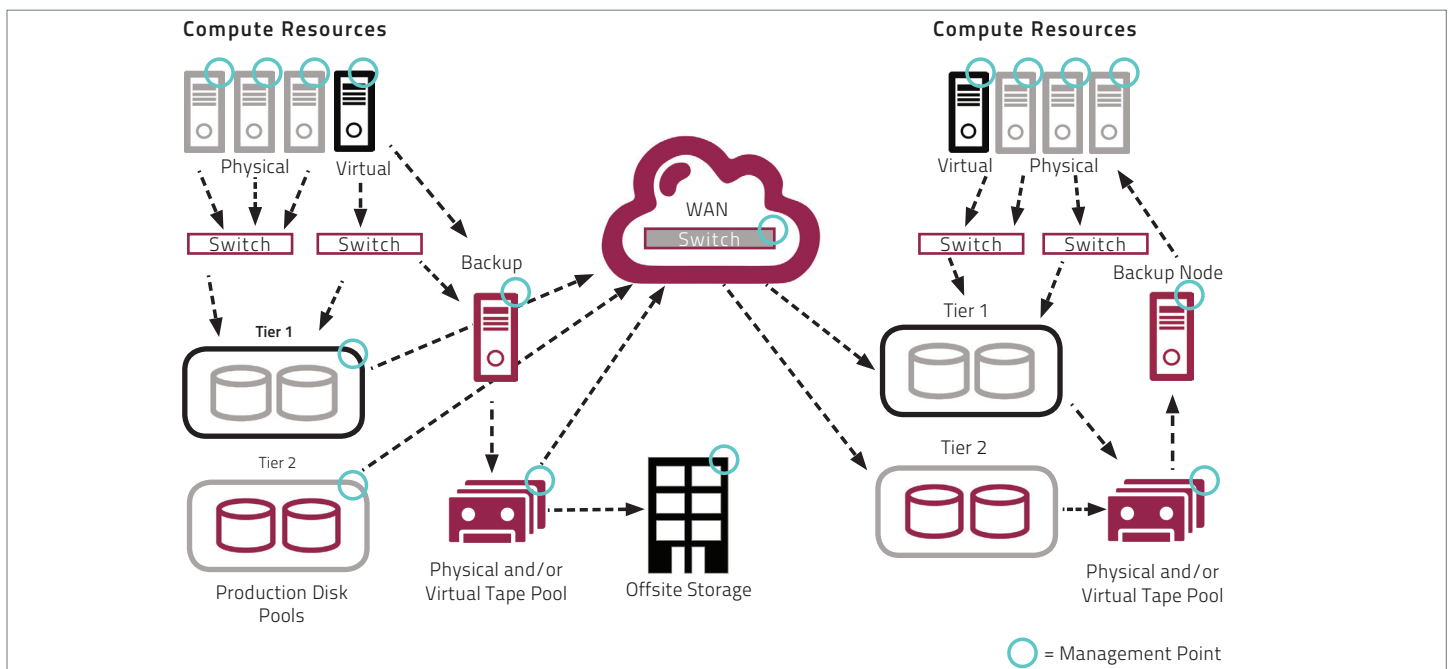


Figure 1: Traditional architecture

The need to manage multiple individual elements increases cost and complexity. Every circle depicts a management point under this model. Every management point adds complexity and operational cost.

Some common aspects of the traditional IT model are:

- Tape is used as primary backup medium.
- Disaster recovery testing is difficult and time consuming. Environment changes need to be captured and documented. Manual processes need to be changed and tested accordingly.
- Batch processing is used to streamline tasks, but they are managed individually.
- Operations management is segmented into multiple groups responsible for subsets of the IT environment. (example: network, storage, server management, database management).
- Disaster recovery planning for structured and un-structured data typically involves shipping tapes offsite.
- Storage Array-based replication is used to move critical data over the network to a similar storage array at the disaster recovery site.
- Virtual servers may be introduced to reduce server footprint, but a physical server plant still exists for high performance applications.
- The disaster recovery site needs to mirror the primary site for most mission-critical applications, including similar servers and storage resources.
- Costs for Tier 1 storage are typically high based on the requirement for array-based licenses to provide advanced data services such as snapshots and replication.
- Some virtual services, such as virtual tape, may be present to alleviate backup window issues.
- Although a proven and reliable architecture, the traditional model has numerous crucial shortcomings, particularly with regards to data security, storage efficiency, and total cost of ownership (TCO).

TRADITIONAL VS. OPTIMIZED ARCHITECTURE

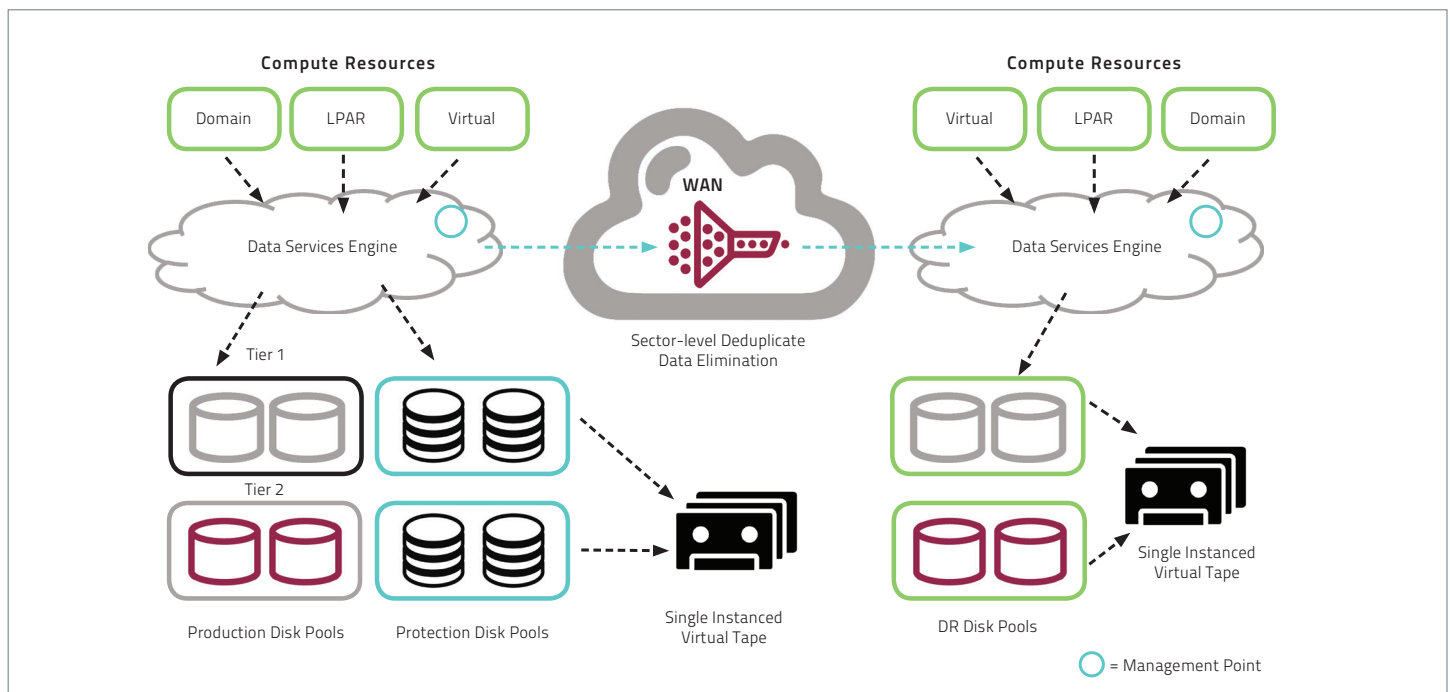


Figure 2: Optimized architecture

ODS MODEL ARCHITECTURE

The Optimized Data Services model proposes an architecture which leverages intelligent abstraction to virtualize the infrastructure and automate most of the mundane tasks required for day-to-day operations. Disk rather than tape is used for protection and recovery. Storage is optimized and the number of data copies are reduced by thin provisioning, data deduplication, and snapshots. Resiliency and reliability of data access is improved via virtual storage pooling and automated information lifecycle management (ILM).

The optimized model automates data protection services by implementing continuous protection rather than a once a day batch-job based backup process. The protection process leaves the primary data untouched to assure high performance. All data writes to disk are continually protected and stored on a separate pool of storage to protect applications from primary disk failure. A data journal is stored locally and remotely, along with snapshot copies taken at regular intervals in coordination with the application, to achieve rapid recovery from disk to any point in time. Data is also deduped, encrypted, and replicated or mirrored to protect from site-level failures. Both locations can be active and productive, or the DR site is located in the cloud. Data retention for near-line access and compliance is provided via a searchable, deduped single instance repository, which is also replicated. The single instance repository provides protection and WAN optimized off-site replication and recovery for non-mission-critical applications. Tape drives are used strictly for deep archives, and created offsite from production data. Tape copies for archives are created automatically based on automation policies created within the intelligent abstraction layer.

All application related protection and recovery occurs from high speed random access disk.

Some common aspects of the ODS Model are:

- Server and storage virtualization are prevalent, enabling intelligent abstraction from physical resources, which commoditizes the IT infrastructure.
- Data deduplication, single instancing, and delta versioning reduce footprint and network bandwidth requirements.
- Intelligent abstraction is implemented at the fabric layer to eliminate the need for intelligence within the storage infrastructure, eliminating the need for storage array licenses.
- Fit-for-purpose components should be used per application service level agreement (SLA).
- Fibre Channel over Ethernet (FCoE) can be used by many applications to help reduce costs.
- SAN based SATA and/or SAS disk types can be used for lower storage tiers. Solid-state disk is implemented as a high performance tier at the abstraction layer, and enabled across all applications.
- Capacity is managed and automatically provided through the intelligent abstraction layer as a capacity-on-demand service to simplify operations and reduce internal hands-on labor costs.
- Storage management tools are enhanced to automate other administrative tasks and support rapid problem resolution.
- A common provisioning model is used to serve all storage users using a tiered method of protection to simplify operations and enhance recovery across all applications.
- User self-serve recovery is enabled at the edge.

BUSINESS OBJECTIVES

The ODS model takes a holistic view of applications and defines an architecture that unifies the numerous and diverse requirement streams within an organization. A key objective is to reduce unit lifecycle cost (as an example, the cost per terabyte of storage) while achieving world-class performance and reliability through several means:

Intelligent Abstraction through virtualization. A key aspect of the Optimized Data Services approach is to apply intelligent virtual abstraction in order to simplify data and application movement and reduce complexity. Once virtualized under a single platform, it becomes much easier to create automation policies that enforce specific service levels for explicit or pooled data sets. Intelligent abstraction simplifies the grouping of application data elements for consistency and recovery purposes. Applications and data are no longer hampered by physical constraints, or vendor intricacies. (i.e., Storage LUNS which require like to like storage for offsite replication). Intelligent abstraction makes IT simple, and becomes the foundation for the next generation software defined datacenter.

Data retention for regulatory compliance. Regulations such as SEC Rule 17a-4 and HIPAA create legislative mandates for the long-term retention of certain segments of data. Purchasing choices must be made carefully to ensure participating vendors conform to and support regulatory compliance, and are aligned with business requirements and not personal or political relationships.

Curtailement of data growth. Growing volumes of stored data is a natural result of business activity. The ODS model proposes methods to achieve significant reductions in the cost per unit of storage to manage data growth. Examples of such methods include the reduction of duplicate data on primary and backup storage elements including WAN requirements for data in flight, and deduplication for backup and archive data. Leveraging deduplication across backup data can provide immediate business savings in terms of reduced tape media and management costs. Solutions to reduce non-business use of storage can also be offered progressively. Data management tools are introduced to find opportunities to delete data which is no longer relevant to the organization, or data is consolidated and mined to increase its value.

Other areas. The optimized data services architecture also focuses on development and refinement across the following areas. By approaching these areas with an optimized, service-oriented model in mind, organizations can improve efficiencies in the unit cost of managed storage:

Utilization. Increase utilization levels and improve processes for allocation of storage volumes.

Automation. Automate technical activities and enable administrative staff to manage larger storage volumes, provide end users with tools to manage their own storage.

Backup. Backup is the most fragile and costly component of storage services. A re-think of the overall technology and operations is necessary. Cloud should be considered.

Procurement. Storage procurement costs should decrease, not increase. Storage hardware should be acquired in a "just-in-time" manner.

Operations. Islands of storage that are managed separately give rise to duplication of infrastructures, the inability to share resources, and different solutions to fix the same issues. The IT infrastructure must be viewed and treated as an integrated, unified entity.

Business applications. Storage costs and utilization levels need to be gauged in order for organizations to better manage the consumption and protection of business applications.

Migration. Migration of storage environments is an ongoing activity that must be streamlined. Migration of data volumes (user or application moves) must be facilitated to any group location or, where appropriate, to an external location or provider.

STRATEGIC PLANNING

CIOs and data center managers should review their current infrastructure and total cost of ownership (TCO) structure in order to map the current model with future requirements and consideration of the new technology landscape. A new “end-state vision” should be developed to provide a roadmap and enable IT to anticipate the challenges of data growth. Consideration should be made to address new technologies and standards to enable cost containment.

Based on recent surveys, cost reductions of up to 70% are attainable through application of the optimized data services model. Benefits are attained only as quickly as the new strategic architecture strategy is deployed and refined, and legacy systems are migrated to it. Once the existing environment is assessed and quantified financially, a phased approach for implementation of the ODS model should be developed to address the most problematic areas first, where return on investment will provide immediate results.

There are several basic questions IT managers should ask their team to determine whether solutions or projects currently under consideration will contribute to an optimized data services model:

1. Does the solution simplify operations?
2. Can we use this solution across all platforms and applications?
3. Does the solution leverage existing assets?
4. Does the solution leverage, enhance, or automate current policies and procedures?
5. Can we implement the solution based on the budgetary savings it provides?
6. Does the solution enhance access and availability to data, improving the RTOs/RPOs for mission-critical applications?

The first step towards implementation begins with creating strategic policies based on data classification. Additional policies can be used within corporate IT environments to minimize the costs associated with data management. Below is a recommended 10-step guideline for cost containment within an optimized IT infrastructure.

1. Classification of data
2. Policy creation based on data classification
3. Consolidation of the infrastructure via intelligent abstraction and modular reference architectures
4. Virtualization and pooling heterogeneous storage and server resources
5. Placing data in correct storage pools (SSD/FC/SATA/NAS/TAPE)
6. Using correct protocols per policy (FCP/FCoE/iSCSI/VI/IP/FC-IP/CIFS/NFS/DAFS)
7. Using the correct topology per class (IB/FC-SW/FCAL/IP/VI/SATA)
8. Use standards-based management (SMI-S/WEBM/CIM/SOAP/XML/ITIL)
9. Automation
10. Documentation

IT INFRASTRUCTURE ASSESSMENT

The goal of an ODS assessment is to find ways to improve operations and service levels in support of critical applications while reducing costs. The following areas must be included as part of the assessment:

- IT operations (provisioning, data mobility, application rollout and support)
- Network overview (LAN and WAN design and costs)
- Backup design and operation (backup and recovery times, retention, archiving, etc.)
- Storage infrastructure
- Disaster recovery capabilities
- Application recovery

The ODS model should be used as a guideline to determine which new technological investments could be applied to improve utilization of current investments. The assessment results should determine where improvements could be made to simplify operations, improve service levels, reduce costs, and provide better means of protection and recovery for critical applications.

Network. The LAN and WAN infrastructure should be robust enough to provide high performance for all client access while providing residual bandwidth for data movement between locations. Since WAN bandwidth is critical for data replication and in turn disaster recovery, it should be robust enough to handle all production workloads while maintaining bandwidth to keep data in sync between sites to meet stringent recovery point objectives (RPO).

Servers and Applications. In the world of open distributed computing, server virtualization enables cost reductions through the ability to use less physical hardware for application services. This in turn enables more cost-effective disaster recovery in that fewer physical servers are required at a remote site. IT should take advantage of running applications on a robust modular, “blade server” GRID, or Cloud based infrastructure. Domain-based and LPAR-based physical servers could be used for high performance applications, or service-based solutions should be procured from an MSP or cloud vendor.

Storage Area Network (SAN). The SAN can be one of the most costly areas of IT, especially as it pertains to annual OpEx, which can be four times more costly than CapEx. All available measures should be taken to provide a robust, high performing and reliable infrastructure for data storage in the most cost-efficient manner as possible. Virtualization should be used to commoditize disk resources, and enable data mobility between storage from any vendor.

Backup Environment. Backup is the biggest headache facing most organizations. It is important for an IT department to take an inventory of the existing backup environment, business objectives for RTO and RPO for applications, and the biggest issues it faces from a backup and recovery perspective. In an optimized IT model, backup and recovery is typically the largest benefactor from a cost savings perspective. Cloud based services should be reviewed to see if there is financial and operational benefit to moving the backup process to a service provider.

Business Continuity/Disaster Recovery (BC/DR). Organizations benefit dramatically from an optimized IT model in terms of reduced BC/DR costs. Data deduplication and delta versioning can be used across all services to reduce storage and WAN costs. Consolidated FCoE or iSCSI connectivity can be used in lieu of a SAN at the disaster recovery site. Fabric-level replication can be used to achieve single console disaster recovery across all storage classes and vendors. Cloud based services should be reviewed to see if there is financial and operational benefit to moving the disaster recovery process to a service provider.

DESIGNING A SERVICE LEVEL AGREEMENT (SLA) FRAMEWORK

The SLA framework of the ODS model provides data storage services using different service levels to meet a diverse set of business requirements at an appropriate cost. In order to do so, the following service levels need to be defined:

- User file services
- Shared file services
- Low-performance application and database storage services
- Medium-performance application and database storage services
- High-performance application and database storage services

Recommended best practices for designing the SLA framework include the following steps:

- Provide storage services at the lowest unit cost. Storage virtualization within the framework is leveraged to commoditize disk.
- Directly address 80% of most storage requirements while being flexible enough to meet the remaining 20% through other service options.
- Include application-focused operational recovery as a standard on all services.
- Include cloud-based backup and recovery as an option on critical services where internal services may be lacking.
- Leverage continuous protection technology for mission-critical applications; include snapshot-based recovery from disk as a primary method for most applications.
- Leverage virtual tape with deduplication integrated with traditional backup as a method for disk-based disaster recovery for non-critical applications.

COST/BENEFIT ANALYSIS

One of the strategic goals of the assessment is to provide a financial analysis of the total cost of ownership (TCO) savings provided by an ODS model implementation.

The following table lists some of the main areas where IT budget is typically spent within an organization. In an optimized IT model, the expenses in maroon are positively impacted, reducing overall costs.

CAPEX AND OPEX: SOLUTION TCO SAVINGS

HARD COSTS	SOFT COSTS
Labor costs per year	Downtime cost per hour
Current hardware purchase cost	Cost of operations
Current software purchase & license cost	Provisioning efficiency: Time to provision
Software maintenance per year	Revenue loss during outage
Hardware maintenance per year	Missed opportunity loss
Floor space \$/sq foot	Customer loss (new customers)
Media costs per year	RTO impact (cost per hour during DR)
Tape storage (local and off-site)	RPO impact (cost per lost transaction)
Tape transport costs per year	Server rebuild time
Network costs (WAN/LAN)	Seat time (employees doing nothing)
Additional component costs (ports, extenders, HBA)	Image damage (customer perception)
DR equipment costs (FC-IP, Decru, Riverbed)	
Environmental (power, cooling, etc.)	
Monitoring costs	

Figure 3: Solution value

CONCLUSION

The optimized data services model involves a paradigm shift in approaches to IT data management and datacenter infrastructure technology that can be summed up as such:

- Operational structure simplification and standardization to improve efficiency of support staff
- Unifying storage and servers into an infrastructure that can be managed by one group
- Improving management processes and tools
- Leveraging intelligent abstraction to provide complete data mobility, automation of processes, and commoditization of infrastructure

When careful, educated choices are made based on the merits of the technology and applicability within the existing environment, with no deference to vendor relationship or political biases, the following benefits can be achieved:

- Dramatic increases in operational efficiency
- Better alignment between the business and IT organizations
- Improved application reliability and recovery
- Reduced risk
- Optimized utilization of existing IT investments
- Improvements in disaster recovery capabilities and costs
- Reduced overall operating expenses

Corporate Headquarters

2 Huntington Quadrangle, Suite 2S01
Melville, NY 11747
Tel: +1.631.777.5188
salesinfo@falconstor.com

European Headquarters

Landsberger Str. 312
80687 Munich, Germany
Tel: +49 (0) 89.41615321.10
salesemea@falconstor.com

Asia Headquarters

Room 1901, PICC Office Tower
No. 2, Jian Guo Men Wai Street
Chaoyang District
Beijing, 100022 China
Tel: +86.10.6530.9505
salesasia@falconstor.com

Information in this document is provided "AS IS" without warranty of any kind, and is subject to change without notice by FalconStor, which assumes no responsibility for any errors or claims herein. Copyright © 2015 FalconStor Software. All rights reserved. FalconStor Software and FalconStor are registered trademarks of FalconStor Software, Inc. in the United States and other countries. All other company and product names contained herein are or may be trademarks of the respective holder. BPOITWP150224