



RAIDIX 5.0.1 PRODUCT FEATURES

Document version 3.0

CONTENTS

What is RAIDIX 5.0.1?	4
How it works	4
Duplication of Hardware Components	5
Network Attached Storage.....	6
Distinctive NAS Features in RAIDIX	6
Active Directory	7
NFS + LDAP Integration	7
Hybrid Storage	7
Clustering System.....	7
RAID Engine	8
RAID Levels.....	8
RAID 6.....	8
RAID 7.3	8
RAID N+M.....	8
RAID 50, 60, and 70.....	8
High-speed reconstruction of RAID 6, RAID 7.3	9
Replication	9
Replication Protocols	9
NVDIMM: Non-Volatile Memory Protecting Write-Back Cache.....	9
SAN Optimizer	9
Bonding of SAN drives in DC.....	10
Automatic Write Through.....	10
Caching Algorithms.....	10
Protection from Silent Data Corruption.....	10
Scheduled Background Integrity Check	10
Partial Reconstruction	11
Advanced Reconstruction.....	11
Masking	11
System Monitoring Module	11

Drive Scan	11
SparePools.....	12
Universal Host Type	12
QoSmic (Quality of Service)	12
Unlimited LUN Size.....	12
System State Notifications	12
Graphical User Interface.....	12
Simple Version Update	12
Product Features	13
Contact Us	14

WHAT IS RAIDIX 5.0.1?

RAIDIX software is a unique development by RAIDIX LLC intended for creating high-performance storage systems using standard hardware components.

RAIDIX is perfectly suitable for tasks with strict requirements for high performance, reliability and operational continuity owing to the unique proprietary mathematical algorithms and parallel calculations usage.

How it works

RAIDIX 5.0.1 supports both the single-controller mode (one node is used) (Figure 1) and the dual-controller mode (Figure 2), when two nodes are active and have access to the same set of drives. The nodes are hardware-independent components of the storage system, which have their own processors, cache memory, motherboard, power supply units.

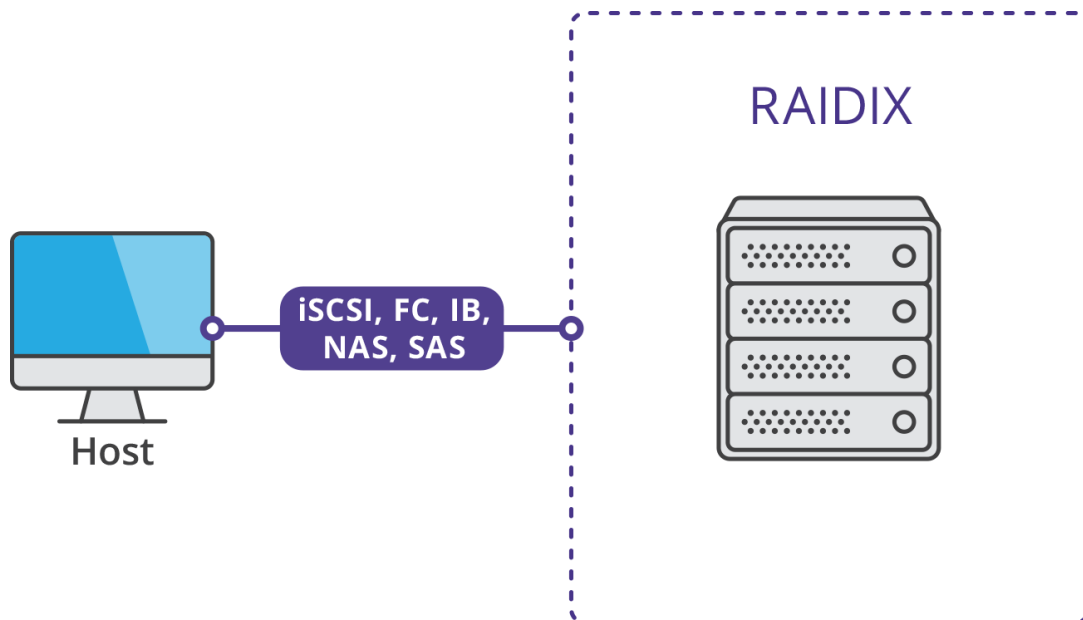


Figure 1. RAIDIX 5.0.1 Single-controller mode

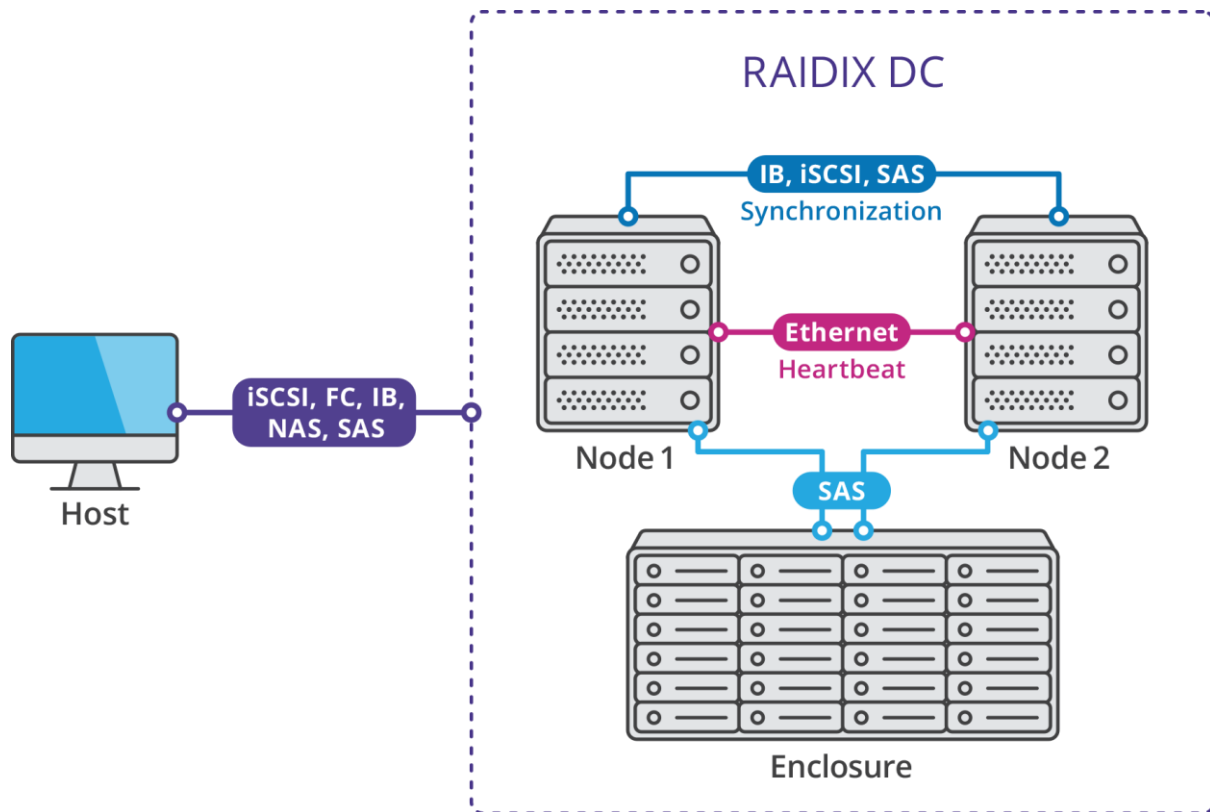


Figure 2. RAIDIX 5.0.1 Dual-controller mode

Duplication of Hardware Components

RAIDIX guarantees continuity of data access and provides fault-tolerance due to:

- duplication of nodes (including motherboards, cache memory modules, power supply units, SAS-controllers, system drives);
- duplication of drive connection channels (both nodes are connected to one set of drives).

Nodes interact over InfiniBand, iSCSI (over Ethernet) and LSI SAS interfaces, enabling to synchronize data and caches state.

Due to the standby cache synchronization in both directions, the remote node always contains an up-to-date copy of data in the local node cache. Thus, if one node fails, the other node transparently takes the entire workload, allowing the administrator to fix appeared errors without stopping the system.

Duplication of hardware components and interfaces provides the protection against the following failures:

- failure of one of the hardware components (CPU, motherboard, power supply unit, controller, system drive);
- enclosure connection interface failure (SAS-cable, I/O-module);
- power-off of one of the nodes;
- occurrence of software errors on one of the nodes.

NETWORK ATTACHED STORAGE

RAIDIX 5.0.1 enables working the Network Attached Storage (NAS). Network system architecture represents a NAS server combined with storage system on RAIDIX platform, which interacts with client computers through SMB/CIFS, NFS, FTP, and AFP protocols (Figure 3).

Main advantages of working with NAS on RAIDIX 5.0.1 platform are:

- lower exploitation cost in comparison to SAN;
- possibility of usage over the local network;
- possibility of file sharing, i.e. simultaneous user access to large amounts of data.

RAIDIX 5.0.1 NAS functionality includes an ability to create and edit shares with adjustable parameters (path, protocol, access permissions, visibility and host selection).

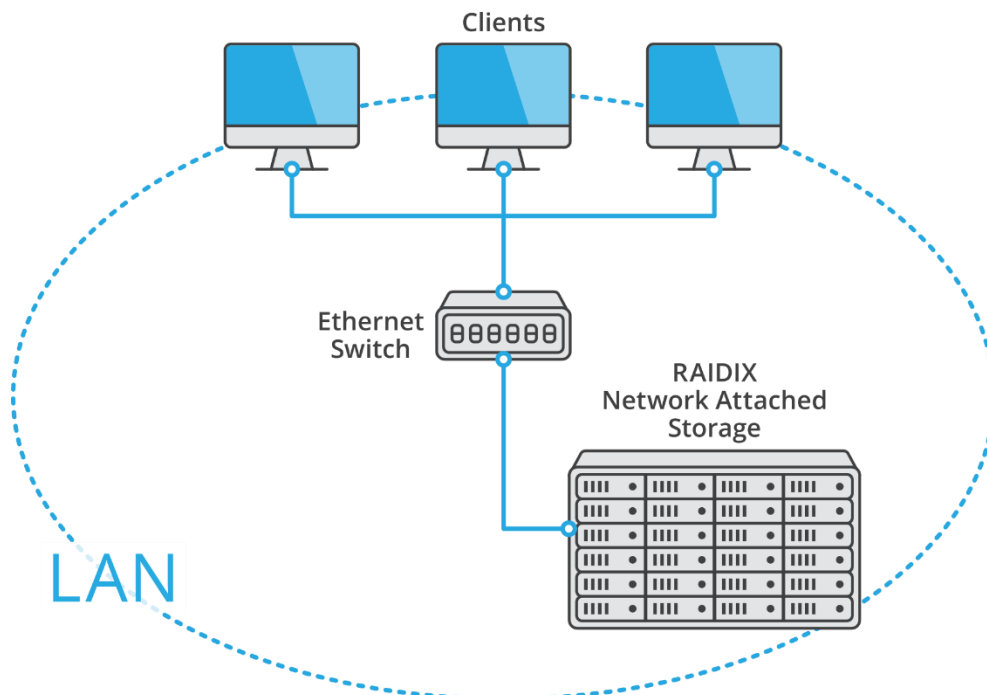


Figure 3. NAS Architecture in RAIDIX

Distinctive NAS Features in RAIDIX

- Clusterization support: NAS is available in dual-controller (Active-Passive) mode.
- WORM (Write once read many) option for SMB shares: ability to prohibit file modification after its creation for a long period (5 years by default).
- Ability to select the type of the file system: XFS or ext4.
- Quoting support: ability to specify the amount of LUN size, available for the particular share users or user groups.

Active Directory

While working with SMB shares, integration with Active Directory is available. Active Directory allows several domain users or group of users to work with shares, assigning different access permissions to them.

NFS + LDAP Integration

In RAIDIX 5.0.1, we implement the ability to connect to LDAP servers and enable LDAP users to work with NFS shares and set access rights to files and directories.

LDAP (Lightweight Directory Access Protocol) is an application protocol intended for accessing the directory information services based on X.500. We implement the ability to connect to an LDAP server and access to users and groups from the server.

HYBRID STORAGE

We implemented the hybrid storage containing both SSDs and HDDs. Use SSDs for caching requests to HDDs. As SSDs, compared to HDDs, are better adapted for random read and write workload, SSD cache provides a higher rate of access to the most often used data ('hot' data) on the hard drives as long as the workload is random.

Clustering System

Clustering system of RAIDIX 5.0.1 generation software allows you to create a fault-tolerant high-performance cluster (by configuring dual-controller mode) and place arrays asymmetrically on the nodes. Each RAID can only be active on one of the nodes through which access to RAID resources is provided.

Implemented system architecture:

- increases fault tolerance due to automatic and manual nodes modes switching (Failover);
- helps to increase system performance due to ability to perform RAID migration from any node of the cluster to balance the nodes' workload. During the RAID migration the array preference parameter (Affinity) changes, RAID becomes active on the other node.

RAIDIX clustering system (Cluster-in-the-Box) provides high availability of services and allows system administrators:

- to install additional software that supports clusterization¹;
- to activate or deactivate services on the nodes in accordance with cluster events.

¹ The functionality is available for partners who create their own end-user products and solutions.

RAID ENGINE

RAID Engine – a RAID configuration that optimizes read/write speed for the different drive types. The RAIDIX 5.0.1 software provides following engines:

- Generic – a configuration, at which both HDDs and SSDs can be used in a RAID. Generic RAID is a classical RAIDIX RAID with high level of sustainable performance and data availability.
- ERA_Flash – an optimized for SSDs (NVMe, SAS, and SATA) configuring that uses advantages of flash devices.

RAID LEVELS

RAIDIX 5.0.1 software allows creating the following RAID levels: 0, 1, 5, 6, 7.3, 10, N+M, 50, 60, and 70.

RAID 6

RAID 6 is the level of interleaving blocks with double parity distribution, based on proprietary RAIDIX mathematical algorithms. RAID 6 shows improved performance, since each drive processes the I/O requests (entries) independently, allowing parallel access to the data. RAID 6 can sustain complete failure of two drives in the same group.

RAID 7.3

RAID 7.3 is the level of interleaving blocks with triple parity distribution, allowing restoring data in case of a failure of up to 3 drives. RAID 7.3, based on proprietary high-performance RAIDIX algorithm, allows achieving high levels of performance without an additional load on the CPU.

RAID 7.3 has a higher degree of reliability than RAID 6 for three checksums are calculated using different algorithms. The capacity of 3 drives is allocated for checksums.

RAID 7.3 is highly recommended for RAIDs larger than 32 TB.

RAID N+M

RAID N+M is the level of interleaving blocks with M checksums, based on proprietary RAIDIX mathematical algorithm. RAID N+M allows user to choose the number of drives for checksums allocation. RAID sustains complete failure up to 32 drives in the same group (depending on the number of parity drives).

RAID 50, 60, and 70

RAID 50, 60, or 70 – RAIDs 0, the segments of which are RAIDs level 5, 6, or 7.3 (called “groups”) accordingly. Such a combination may show better performance with reduced latency. While the size of one, two, or three drives accordingly in each group are allocated to parity data, these RAIDs are recoverable from 1, 2, or 3 failures in each stripe.

High-speed reconstruction of RAID 6, RAID 7.3

RAIDIX 5.0.1 implements a high-speed reconstruction of RAID 6, RAID 7.3 that *runs up to 6 times faster* in comparison to similar hardware controllers and without performance degradation.

Reconstruction runs in background mode during drive replacement and does not affect the client.

REPLICATION

RAIDIX 5.0.1 supports the replication feature. Replication is data duplication on two storage systems.

Replication meets the following requirements:

- Business continuity;
- Increased data storage reliability;
- High data integrity without performance degradation.

Replication Protocols

In RAIDIX 5.0.1, the following replication protocols are available:

- *Synchronous* – initiator receives write confirmation only after data is safely replicated to the secondary LUN. Use of synchronous protocol ensures data integrity but may lead to degradation of performance. This type of replication is demanding in terms of channel capacity and network latency.
- *Asynchronous* – write is acknowledged on the initiator after data is written on the primary LUN. The primary storage tracks data storage and replicates in to the secondary LUN periodically. Asynchronous protocol help to avoid performance deterioration with replication enabled.

NVDIMM: NON-VOLATILE MEMORY PROTECTING WRITE-BACK CACHE

We implemented write-back cache protection by using non-volatile memory NVDIMM-N, which prevents data loss even in case of power failure on the node without performance degradation.

Persistent memory combines the benefits of conventional storage devices and high processing capacity of DRAM. The peculiarities of persistent memory are byte addressing with its high data rates and very small delays.

SAN OPTIMIZER

SAN Optimizer is the technology of virtualization, which enables to organize several independent storages into a joint virtual storage. SAN optimizer allows to connect targets on external storages to RAIDIX and to work with them in the same way as with local drives; for instance, to create RAIDs and LUNs. SAN Optimizer enables virtualization (forwarding) of a drive as a block device without being modified.

SAN Optimizer has the following advantages:

- High performance due to caching in RAIDIX;
- Enhanced reliability due to creation of RAIDs;
- Budget solution of data storage, no additional purchases of expensive equipment needed.

Bonding of SAN drives in DC

While working in DC, it is possible to add SAN drives only on one node. Bonding is intended to provide access to these drives on the second node.

AUTOMATIC WRITE THROUGH

We implemented the automatic enabling of Write Through mechanism for sequential write. Automatic Write Through may improve performance for mixed write pattern in DC mode in case synchronization is a bottleneck.

The feature automatically selects between Write Back or Write Through modes on the basis of information from sequence detector. If the write pattern is determined as random, the data is written to the cache, i.e., Write Back mechanism is used. If the write pattern is determined as sequential, the data is written directly to the drives, i.e. Write Through mechanism is used.

The administrator have an ability to turn on/off write through option manually.

CACHING ALGORITHMS

RAIDIX 5.0.1 supports effective replacement algorithms: *Cycle*, which is adapted for the random write workload, and *LRU* algorithm, which is optimal for the sequential write workload.

PROTECTION FROM SILENT DATA CORRUPTION

Silent Data Corruption may be caused by failures in drivers and drive firmware, memory errors, drive-head crashes and similar software and hardware problems. Silent errors occurred during write operations to the drive are the most hazardous, as they go undetected by drive firmware, host operating systems, and may result in the corruption of data structure leading to the subsequent loss of data.

RAIDIX team implemented the forward silent error correction algorithm that analyzes RAID metadata to detect and fix corruption while regular drive operations are performed, without performance degradation.

SCHEDULED BACKGROUND INTEGRITY CHECK

The check is suitable for:

- correcting of SDC errors;
- excluding of failed drives;
- reducing corruption probability of infrequently used data.

RAID integrity check starts automatically in accordance with the schedule set by administrator. Integrity check runs in the background and doesn't block the normal work of the system.

PARTIAL RECONSTRUCTION

Hard drives capacity growth and therefore slower drive recovery time lead to the increase of probability of drives failure and data loss risk.

RAIDIX 5.0.1 has a mechanism of RAID *Partial Reconstruction* that allows restoring only a particular area containing corrupted data on a hard drive, reducing the array recovery time. Partial reconstruction is particularly effective for large arrays.

ADVANCED RECONSTRUCTION

RAIDIX 5.0.1 provides the improved mechanisms of *Advanced Reconstruction*, which enables to optimize performance of read operations. The effect is achieved by eliminating the drives with the lowest read rates.

MASKING

In terms of RAIDIX software, masking is a process of determining the level of host LUN access.

Target Masking Rules allow system administrators to specify an iSCSI/SAS/Fibre Channel/InfiniBand/12G SAS target for a LUN; the LUN will be available to host through this target.

Host Masking Rules allow to specify the host access to a LUN or manage host access to all LUNs simultaneously.

SYSTEM MONITORING MODULE

System monitoring section assists a user to gather information about the controller and enclosure sensors; about the storage system's performance; about presence/absence of the system faults and provides an ability to download the system logs.

Performance monitoring module helps to map up configuration, estimate and increase the performance both inside of the storage system itself and during the data transfer process.

DRIVE SCAN

RAIDIX 5.0.1 supports the functionality to scan drives of an array on performing read/ write operations. Scan results show the number of read/write commands performed by the system at different time intervals. Analysis of the results allows indicating drives with low performance. The results are presented in a table.

SPAREPOOLS

In RAIDIX 5.0.1, *SparePool* functionality is implemented. Administrator is allowed to create one or more SparePools and add one or several drives in a SparePool; one SparePool can be assigned to one or more RAIDs that administrator wants to protect with hot spare.

UNIVERSAL HOST TYPE

In RAIDIX 5.0.1, the system automatically specifies the default host type, providing LUN interaction with hosts of any supported operating systems. The functionality helps to reduce the probability of errors during the system configuration and connection of hosts.

QOSMIC (QUALITY OF SERVICE)

QoSmic service implemented in RAIDIX 5.0.1 allows user to assign different priority for hosts in the system (Realtime parameter) according to the application type (important or not) running in the host.

UNLIMITED LUN SIZE

Extend LUNs on several RAIDs. LUN size is restricted only by the total space on all RAIDs created in the system. LUN can take fully or partially any number of RAIDs.

SYSTEM STATE NOTIFICATIONS

RAIDIX 5.0.1 supports sending notifications on the system state messages via SMTP and SNMP protocols. The notifications can contain information on the following subjects: RAID, iSCSI, Drive, Network, SparePool, License, Controller Sensors, and Enclosure Sensors. The user can select the notification message types: Info, Warning, or Error.

GRAPHICAL USER INTERFACE

The features of RAIDIX 5.0.1 GUI are:

- 5 supported languages: English, Chinese, Korean, Russian, Japanese;
- HTTPS encryption,
- Support of wizards, enabling fast execution of basic operations;
- Access Protection (user authorization);
- Independent of the operating system.

SIMPLE VERSION UPDATE

RAIDIX offers the easiest way to update the system version with RAIDIX GUI.

PRODUCT FEATURES

Drives Supported	
	<ul style="list-style-type: none"> • 3.5" SAS / NL-SAS / SATA HDDs • 2.5" SAS / NL-SAS / SATA HDDs • SATA / SAS SSDs • NVMe
Operational Features	
RAID levels	<ul style="list-style-type: none"> • Generic RAID: 0, 1, 5, 6, 7.3, 10, and N+M; • ERA RAID: 0, 1, 5, 6, 7.3, 10, 50, 60, and 70.
Max number of drives in one RAID	64.
Max number of drives in the system	600.
Max LUN size	Limited only by total free space of all RAIDs created in the system.
Max LUNs	447.
Heterogeneous System Architecture Support	LUNs could be accessed simultaneously through all available system interfaces.
iSCSI	MPIO, ACLs, CHAP authorization, LUN masking, CRC Digest.
Sessions supported	1024.
Maximum number of directly connected hosts (depends on the hardware platform)	32.
Client OS support	<ul style="list-style-type: none"> • Mac OS X 10.6.8 and older, 10.7, 10.8, 10.9, 10.10; • Microsoft Windows Server 2008, 2008 R2, 2012, 2012 R2, 2016, 2019; • Microsoft Windows 8 (32-bit, 64-bit); • Red Hat Linux, SuSE, ALT Linux, CentOS Linux, Ubuntu Linux.
Virtualization platforms support	<ul style="list-style-type: none"> • VMware ESXi 5.5, 6.0, 6.5, 6.7; • KVM (Kernel-based Virtual Machine); • RHEV (Red Hat Enterprise Virtualization); • Microsoft Hyper-V Server, XenServer, Proxmox VE.
High-speed communication channels support	Fibre Channel; InfiniBand (FDR, QDR, DDR, EDR); iSCSI; 12G SAS.
MPIO	<ul style="list-style-type: none"> • Microsoft ® Windows® Server 2008, 2008 R2, 2012, 2012 R2, 2016; • VMware ESXi 5.5, 6.0, 6.5, 6.7; • Linux: (including but not limited to) Red Hat Linux, SuSE, ALT Linux, CentOS Linux, Ubuntu Linux.
NAS protocols support	SMB/CIFS, NFS, FTP, AFP.
Applications support	StorNext, MetaSAN, Lustre, FhGFS, GPFS, CXFS, Hyper FS.

Tested with	AJA, Matrox, Final Cut Pro, Blue Fish, BlackMagic, Iridas, VeeamBackUP, NetBackup, FalconStore, Adobe Premier.
Administration and Management	
UI	Multi-language intuitive management Web-interface.
Authorization	User name and password.
Remote access	HTTP, HTTPS, SSH.
Information on sensors	Controller and enclosure.
Notifications	E-mail, UI hardware failure notification.
Others	
Warranty	Optional extended warranty.
Support	One year of basic technical support service is included with license purchase. Optional extended support is available.
Software Update	Free within the technical support period, or can be purchased separately beyond the support period.

CONTACT US

To get in touch, please use the appropriate links below.

- Learn more: www.raidix.com
- E-mail: request@raidix.com
- Call us: +7 (812) 622 16 80.