

Jini Network Technology PROJECT REPORT

MASTER OF COMPUTER APPLICATIONS BY AMIT KUMAR SINHA

NATIONAL INSTITUTE OF TECHNOLOGY CALICUT
NIT CAMPUS P.O, KOZHIKODE 673601, KERALA

DEPARTMENT OF COMPUTER ENGINEERING

National Institute of Technology Calicut

Department of Computer Engineering CERTIFICATE This is to certify that the work reported in this project report entitled "Jini Network Technology" is a bonafide record of the work done by AMIT KUMAR SINHA (Y2M001), a student in the Department of Computer Engineering, National Institute of Technology Calicut .

Co-ordinator:

Head of the Department:

Acknowledgement

I would like to put on records my sincere thanks to: Dr V.k. Govindan Head of Computer science and Engineering Department, my guide for this seminar, who helped me in preparing for this seminar. Mrs. Priya Chandran and Miss Nisha K.K, for all guidance I have received. My friends who's wholehearted support helped me at all stages of the preparation of this seminar.

Table of Contents	1	Introduction-----
1.1 What is Jini?	2	-----
1.2 Goals of the System-----		
1.3 How Jini Works-----	3	
1.4 Environmental Assumptions -----	4	
	5	
2. System Overview		
2.1 Key Concepts-----		
2.1.1 2.1.2 2.1.3 2.1.4 2.1.5 2.1.6 2.2 Services-----	6	Java remote Method Invocation-----
		6 Security-----
		7 Leasing-----
2.2 Services-----	7	7 Transaction-----
Events-----		
	8	
Component Overview -----		
	8	
2.2.1 Infrastructure-----		
2.2.2 Programming Model-----	9	
2.3 Service Architecture-----	10	
	12	
2.3.1 Discovery Process-----		
2.3.2 Join Process-----	12	
2.1.3 Lookup Process-----	13	
2.4 Separation of Interface and implementation-----	14	
	15	
3. Performance Analysis-----		17
Conclusion-----		18
4. References-----		1
5. References-----		
	8	

ABSTRACT Jini is a network architecture for the construction of the distributed systems where scale, rate of change and the complexity of the interactions within and between networks are extremely important and cannot be satisfactorily addressed by existing technology. This technology provides a flexible infrastructure for delivering services in a network and for creating spontaneous interactions between clients that use these services regardless of their hardware or software implementations. Concepts introduced by this technology such as leasing discovery and the movement of objects over a network vastly simplifies interaction on a network embedding a great degree of resiliency and scalability within a distributed system. Companies are using this technology to build communities of services out of simple piece of hardware and software that have been built according to it's specification. Examples ranges from business enterprises, medical systems, financial applications, and battlefield implementations, to a Jini enabled car or a Smart House. It's source code is available at no cost. There are no fees of any kind associated with the licensing and use of this technology. The Jini architecture is the only technology in the marketplace that moves Objects around the network and relegates wire protocols to mere implementation details. It's vision is "network anything , anytime, anywhere !".Due to the simplicity of Jini architecture it clearly wins and will be.

1. Introduction:

Under the leadership of Bill Joy and Jim Waldo on January 25, 1999 Sun Microsystems launched Jini technology."Jini isn't an acronym though, it has been stated that it does function as an anti acronym "Jini Is Not Initials". Over the last quarter century, network technology has evolved rapidly and in some unexpected ways. Client/server and multi-tier models operating within a single business enterprise have given way to an Internet/Web environment where services are provided by nodes scattered over a far-flung network. Today, the next generation of network interaction is emerging that will place unprecedented demands upon existing network technologies and architectures. For example, participants in one network will need to directly access and use the services provided by participants in another network. It is in this network environment - one of mindnumbing complexity driven by geometric increases in scale, rate of change, and multiplicity of participant interactions - that the simplicity of the Jini architecture clearly wins. Traditionally, operating systems have been designed with the assumption that a computer will have a processor, some memory, and a disk. When we boot a computer, the first thing it does is look for a disk. If it doesn't find a disk, it can't function as a computer. Increasingly, however, computers are appearing in a different guise: as embedded devices that have a processor, some memory, and a network connection -- but no disk. The first thing a cellphone does when we boot it up, for example, is look for the telephone network. If it doesn't find the network, it can't function as a cellphone. This trend in the hardware environment, from disk-centric to network-centric, will affect how we organize our software -- and that's where Jini comes in. Jini is an attempt to rethink computer architecture, given the rising importance of the network and the proliferation of processors in devices that have no disk drive. These devices, which will come from many different vendors, will need to interact over a network. The network itself will be very dynamic -- devices and services will be added and removed regularly. Jini provides mechanisms to enable smooth adding, removal, and finding of devices and services on the

network. In addition, Jini provides a programming model that makes it easier for programmers to get their devices talking to each other. Building on top of Java, object serialization, and RMI, which enable objects to move around the network from virtual machine to virtual machine, Jini attempts to extend the benefits of object-oriented programming to the network. Instead of requiring device vendors to agree on the network protocols through which their devices can interact, Jini enables the devices to talk to each other through interfaces to objects.

1.1 What is Jini?

Jini is a set of APIs and network protocols that can help us build and deploy distributed systems that are organized as federations of services. A service can be anything that sits on the network and is ready to perform a useful function. Hardware devices, software, communications channels -- even human users themselves -- can be services. A Jini-enabled disk drive, for example, could offer a "storage" service. A Jini-enabled printer could offer a "printing" service. A federation of services, then, is a set of services, currently available on the network, that a client (meaning a program, service, or user) can bring together to help it accomplish some goal.

Fig.1 Federation of Service available on the Jini network. The idea behind the word federation is based on the Jini view of the network -- there isn't a central controlling authority. Because no one service is in charge, the set of all services available on the network form a federation -- a group composed of equal peers. Instead of a central authority, Jini's runtime infrastructure merely provides a way for clients and services to find each other (via a lookup service, which

stores a directory of currently available services). After services locate each other, they are on their own. The client and its enlisted services perform their task independently of the Jini runtime infrastructure. If the Jini lookup service crashes, any distributed systems brought together via the lookup service before it crashed can continue their work. Jini even includes a network protocol that clients can use to find services in the absence of a lookup service.

1.2 Goals of the System

The overall goal is to turn the network into a flexible, easily administered tool with which human and computational clients can find resources. Resources can be implemented as either hardware devices, software programs, or a combination of the two. The focus of the system is to make the network a more dynamic entity that better reflects the dynamic nature of the workgroup by enabling the ability to add and delete services flexibly. A Jini system consists of the following parts:

• • •

A set of components that provides an infrastructure for federating services in a distributed system A programming model that supports and encourages the production of reliable distributed services Services that can be made part of a federated Jini system and that offer functionality to any other member of the federation

Although these pieces are separable and distinct, they are interrelated, which can blur the distinction in practice. The components that make up the Jini technology infrastructure make use of the Jini technology-programming model; services that reside within the infrastructure also use that model; and the programming model is well supported by components in the infrastructure. The end goals of the system span a number of different audiences; these goals include the following:

•

Enabling users to share services and resources over a network

•
•
•
Providing users easy access to resources anywhere on the network while allowing the network location of the user to change Simplifying the task of building, maintaining, and altering a network of devices, software, and users

The Jini system extends the Java application environment from a single virtual machine to a network of machines. The Java application environment provides a good computing platform for distributed computing because both code and data can move from machine to machine. The environment has built-in security that allows the confidence to run code downloaded from another machine. Strong typing in the Java application environment enables identifying the class of an object to be run on a virtual machine even when the object did not originate on that machine. The result is a system in which the network supports a fluid configuration of objects that can move from place to place as needed and can call any part of the network to perform operations. The Jini architecture exploits these characteristics of the Java application environment to simplify the construction of a distributed system. The Jini architecture adds mechanisms that allow fluidity of all components in a distributed system, extending the easy movement of objects to the entire networked system. The Jini technology infrastructure provides mechanisms for devices, services, and users to join and detach from a network. Joining and leaving a Jini system are easy and natural, often automatic, occurrences. Jini systems are far more dynamic than is currently possible in networked groups where configuring a network is a centralized function done by hand.

1.3 How Jini works Jini defines a runtime infrastructure that resides on the network and provides mechanisms that enable you to add, remove, locate, and access services. The runtime infrastructure resides on the network in three places: in lookup services that sit on the network; in the service providers (such as Jini-enabled devices); and in clients. Lookup

services are the central organizing mechanism for Jini-based systems. When new services become available on the network, they register themselves with a lookup service. When clients wish to locate a service to assist with some task, they consult a lookup service. The runtime infrastructure uses one network-level protocol, called discovery, and two object-level protocols, called join and lookup. Discovery enables clients and services to locate lookup services. Join enables a service to register itself in a lookup service. Lookup enables a client to query a lookup service for services that can help the client accomplish its goals.

1.4 Environmental Assumptions The Jini system federates computers and computing devices into what appears to the user as a single system. It relies on the existence of a network of reasonable speed connecting those computers and devices. Some devices require much higher bandwidth and others can do with much less displays and printers are examples of extreme points. We assume that the latency of the network is reasonable. We assume that each Jini technology-enabled device has some memory and processing power. Devices without processing power or memory may be connected to a Jini system, but those devices are controlled by another piece of hardware and/or software that presents the device to the Jini system and it self contains both processing power and memory. Architectures for devices not equipped with a Java virtual machine 1 (JVM) are explored more fully in the Jini Device Architecture Specification. The Jini architecture gains much of its simplicity from assuming that the Java programming language is the implementation language for components. The ability to dynamically download and run code is central to a number of the features of the Jini architecture. However, the Java technology-centered nature of the Jini architecture depends on the Java application environment rather than on the Java programming language. Any programming language can be supported by a Jini system if it has a compiler that produces compliant bytecodes for the Java programming language.

2 System Overview

2.1 Key Concepts

The purpose of the Jini architecture is to federate groups of devices and software components into a single, dynamic distributed system. The resulting federation provides the simplicity of access, ease of administration, and support for sharing that are provided by a large monolithic system while retaining the flexibility, uniform response, and control provided by a personal computer or workstation. The architecture of a single Jini system is targeted to the workgroup. Members of the federation are assumed to agree on basic notions of trust, administration, identification, and policy. It is possible to federate Jini systems themselves for larger organizations.

2.1.1 Services

The most important concept within the Jini architecture is that of a service as a forementioned. Members of a Jini system federate to share access to services. Jini systems provide mechanisms for service construction, lookup, communication, and use in a distributed system. Examples of services include: devices such as printers, displays, or disks; software such as applications or utilities; information such as databases and files; and users of the system. Services in a Jini system communicate with each other by using a service protocol, which is a set of interfaces written in the Java programming language. The set of such protocols is open ended. The base Jini system defines a small number of such protocols that define critical service interactions.

2.1.2 Java Remote Method Invocation (RMI)

Communication between services can be accomplished using Java Remote Method Invocation (RMI). The infrastructure to support communication between services is not itself a service that is discovered and used but is, rather, a part of the Jini

technology infrastructure. RMI provides mechanisms to find, activate, and garbage collect object groups. Fundamentally, RMI is a Java programming language-enabled extension to traditional remote procedure call mechanisms. RMI allows not only data to be passed from object to object around the network but full objects, including code. Much of the simplicity of the Jini system is enabled by this ability to move code around the network in a form that is encapsulated as an object.

2.1.3 Security The design of the security model for Jini technology is built on the twin notions of a principal and an access control list. Jini services are accessed on behalf of some entity the principal which generally traces back to a particular user of the system. Services themselves may request access to other services based on the identity of the object that implements the service. Whether access to a service is allowed depends on the contents of an access control list that is associated with the object.

2.1.4 Leasing Access to many of the services in the Jini system environment is lease based. A lease is a grant of guaranteed access over a time period. Each lease is negotiated between the user of the service and the provider of the service as part of the service protocol: A service is requested for some period; access is granted for some period, presumably taking the request period into account. If a lease is not renewed before it is freed--either because the resource is no longer needed, the client or network fails, or the lease is not permitted to be renewed--then both the user and the provider of the resource may conclude that the resource can be freed. Leases are either exclusive or non-exclusive. Exclusive leases ensure that no one else may take a lease on the resource during the period of the lease; non-exclusive leases allow multiple users to share a resource.

2.1.5 Transactions A series of operations, either within a single service or spanning multiple services, can be wrapped in a transaction. The Jini transaction interfaces supply a service protocol needed to coordinate a two-phase commit. How transactions are implemented--and indeed, the very semantics of the notion of a transaction--is left up to the service using those interfaces.

2.1.6 Events The Jini architecture supports distributed events. An object may allow other objects to register interest in events in the object and receive a notification of the occurrence of such an event. This enables distributed event-based programs to be written with a variety of reliability and scalability guarantees.

2.2. Component Overview

The components of the Jini system can be segmented into three categories: infrastructure, programming model, and services. The infrastructure is the set of components that enables building a federated Jini system, while the services are the entities within the federation. The programming model is a set of interfaces that enables the construction of reliable services, including those that are part of the infrastructure and those that join into the federation. These three categories, though distinct and separable, are entangled to such an extent that the distinction between them can seem blurred. Moreover, it is possible to build systems that have some of the functionality of the Jini system with variants on the categories or without all three of them. But a Jini system gains its full power because it is a system built with the particular infrastructure and programming models described, based on the notion of a service. Decoupling the segments within the architecture allows legacy code to be changed minimally to take part in a Jini system. Nevertheless, the full power of a Jini system will be available only to new services that are constructed using the integrated model. A Jini system can be seen as a network extension of the infrastructure, programming model, and services that made Java technology successful in the single-machine case. These

categories along with the corresponding components in the familiar Java application environment are shown in Figure 2.1:

Infrastructure Java VM Base RMI Java Java Security
Programming Model Java APIs JavaBeans ... JNDI
Services
Enterprise Beans JTS ...
Java Discovery/Join + Jini Distributed Security Lookup
Leasing Transactions Events
Printing Transaction Manager JavaSpaces Service ...

Figure 2.1: Jini Architecture Segmentation 2.2.1 Infrastructure The Jini technology infrastructure defines the minimal Jini technology core. The infrastructure includes the following:

-
-
-

A distributed security system, integrated into RMI, that extends the Java platform's security model to the world of distributed systems. The discovery and join protocols, service protocols that allow services (both hardware and software) to discover, become part of, and advertise supplied services to the other members of the federation. The lookup service, which serves as a repository of services. Entries in the lookup service are objects written in the Java programming language; these objects can be downloaded as part of a lookup operation and act as local

proxies to the service that placed the code into the lookup service. The discovery and join protocols define the way a service of any kind becomes part of a Jini system; RMI defines the base language within which the Jini technology-enabled services communicate; the distributed security model and its implementation define how entities are identified and how they get the rights to perform actions on their own behalf and on the behalf of others; and the lookup service reflects the current members of the federation and acts as the central marketplace for offering and finding services by members of the federation.

2.2.2 Programming Model

The infrastructure both enables the programming model and makes use of it. Entries in the lookup service are leased, allowing the lookup service to reflect accurately the set of currently available services. When services join or leave a lookup service, events are signaled, and objects that have registered interest in such events get notifications when new services become available or old services cease to be active. The programming model rests on the ability to move code, which is supported by the base infrastructure. Both the infrastructure and the services that use that infrastructure are computational entities that exist in the physical environment of the Jini system. However, services also constitute a set of interfaces, which define communication protocols that can be used by the services and the infrastructure to communicate between themselves. These interfaces, taken together, make up the distributed extension of the standard Java programming language model that constitutes the Jini programming model. Among the interfaces that make up the Jini programming model are the following:

-
-

The leasing interface, which defines a way of allocating and freeing resources using a renewable, duration-based model. The event and notification interfaces, which are an extension of the event model used by JavaBeans components to the distributed environment, enable event -

based communication between Jini technology-enabled services. The transaction interfaces, which enable entities to cooperate in such a way that either all of the changes made to the group occur atomically or none of them occur.

The transaction interfaces introduce a lightweight, object-oriented protocol enabling applications using Jini technology to coordinate state changes. The transaction protocol provides two steps to coordinate the actions of a group of distributed objects. The first step is called the voting phase, in which each object "votes" whether it has completed its portion of the task and is ready to commit any changes it made. In the second step, a coordinator issues a "commit" request to each object. The Jini transaction protocol differs from most transaction interfaces in that it does not assume that the transactions occur in a transaction processing system. Such systems define mechanisms and programming requirements that guarantee the correct implementation of a particular transaction semantics. The Jini transaction protocol takes a more traditional object-oriented view, leaving the correct implementation of the desired transaction semantics up to the implementer of the particular objects that are involved in the transaction. The goal of the transaction protocol is to define the interactions that such objects must have to coordinate such groups of operations. The interfaces that define the Jini programming model are used by the infrastructure components where appropriate and by the initial Jini technology-enabled services. For example, the lookup service makes use of the leasing and event interfaces. Leasing ensures that services registered continue to be available, and events help administrators discover problems and devices that need configuration. The JavaSpaces service, one example of a Jini technology-enabled service, utilizes leasing and events, and also supports the Jini transaction protocol. The transaction manager can be used to coordinate the voting phase of a transaction for those objects that support transaction protocol. The implementation of a service is not required to use the Jini programming model, but such services need to use that model for their interaction with the Jini technology infrastructure. For example, every service interacts with the Jini lookup service by using the programming model; and whether a service offers

resources on a leased basis or not, the service's registration with the lookup service will be leased and will need to be periodically renewed. The binding of the programming model to the services and the infrastructure is what makes such a federation a Jini system not just a collection of services and protocols. The combination of infrastructure, service, and programming model, all designed to work together and constructed by using each other, simplifies the overall system and unifies it in a way that makes it easier to understand.

2.3 Service Architecture: Service from the interactive basis for a Jini system, both at the programming and user interface levels. The details of the service architecture are best understood once the Jini lookup protocols are presented.

2.3.1 the discovery process

Discovery works like this: Imagine We have a Jini-enabled disk drive that offers a persistent storage service. As soon as We connect the drive to the network, it broadcasts a presence announcement by dropping a multicast packet onto a well-known port. Included in the presence announcement is an IP address and port number where the disk drive can be contacted by a lookup service. Lookup services monitor the well-known port for presence announcement packets. When a lookup service receives a presence announcement, it opens and inspects the packet. The packet contains information that enables the lookup service to determine whether or not it should contact the sender of the packet. If so, it contacts the sender directly by making a TCP connection to the IP address and port number extracted from the packet. Using RMI, the lookup service sends an object, called a service registrar, across the network to the originator of the packet. The purpose of the service registrar object is to facilitate further communication with the lookup service.

Fig.2.2 The Discovery Process

By invoking methods on this object, the sender of the announcement packet can perform join and lookup on the lookup service. In the case of the disk drive, the lookup service would make a TCP connection to the disk drive and would send it a service registrar object, through which the disk drive would then register its persistent storage service via the join process.

2.3.2 The join process

Once a service provider has a service registrar object, the end product of discovery, it is ready to do a join -- to become part of the federation of services that are registered in the lookup service.

Fig.2.3 The Join Process 13

To do a join, the service provider invokes the register() method on the service registrar object, passing as a parameter an object called a service item, a bundle of objects that describe the service. The register() method sends a copy of the service item up to the lookup service, where the service item is stored. Once this has completed, the service provider has finished the join process: its service has become registered in the lookup service. The service item is a container for several objects, including an object called a service object, which clients can use to interact with the service. The service item can also include any number of attributes, which can be any object. Some potential attributes are icons, classes that provide GUIs for the service, and objects that give more information about the service. Service objects usually implement one or more interfaces through which clients interact with the service. For example, a lookup service is a Jini service, and its service object is the service registrar. The register() method invoked by service providers during join is declared in the ServiceRegistrar interface, which all service registrar objects implement. Clients and service providers talk to the lookup service through the service registrar object by invoking methods declared in the ServiceRegistrar interface. Likewise, a disk drive would provide a service object that implemented some well-known storage service interface. Clients would look up and interact with the disk drive by this storage service interface.

2.3.3 The lookup process

Once a service has registered with a lookup service via the join process, that service is available for use by clients who query that lookup service. To build a distributed system of services that will work together to perform some task, a client must locate and enlist the help of the individual services. To find a service, clients query lookup services via a process called lookup. To perform a lookup, a client invokes the lookup() method on a service registrar object. (A client, like a service provider, gets a service registrar through the process of discovery, as described earlier in this article.) The client passes as an argument to lookup() a service template, an object that serves as search criteria for the query. The service template can include a reference to an array of Class objects. 14

These Class objects indicate to the lookup service the Java type (or types) of the service object desired by the client. The service template can also include a service ID, which uniquely identifies a service, and attributes, which must exactly match the attributes uploaded by the service provider in the service item. The service template can also contain wildcards for any of these fields. A wildcard in the service ID field, for example, will match any service ID. The lookup() method sends the service template to the lookup service, which performs the query and sends back zero to many matching service objects. The client gets a reference to the matching service objects as the return value of the lookup() method. In the general case, a client looks up a service by Java type, usually an interface. For example, if a client needed to use a printer, it would compose a service template that included a Class object for a wellknown interface to printer services. All printer services would implement this well-known interface. The lookup service would return a service object (or objects) that implemented this interface. Attributes can be included in the service template to narrow the number of matches for such a type-based search. The client would use the printer service by invoking on the service object methods declared in the well-known printer service interface.

3. Separation of interface and implementation

Jini's architecture brings object-oriented programming to the network by enabling network services to take advantage of one of the fundamentals of object-oriented programming: the separation of interface and implementation. For example, a service object can grant clients access to the service in many ways. The object can actually represent the entire service, which is downloaded to the client during lookup and then executed locally. Alternatively, the service object can serve merely as a proxy to a remote server. When the client invokes methods on the service object, it sends the requests across the network to the server, which does the real work. The local service object and a remote server can each do part of the work as well. One important consequence of Jini's architecture is that the network protocol used to communicate between a proxy service object and a remote server does not need to be known to the client. As illustrated in the figure below, the network protocol is part of the service's implementation. This protocol is a private matter decided upon by the developer of the service. The client can communicate with the service

via this private protocol because the service injects some of its own code (the service object) into the client's address space. The injected service object could communicate with the service via RMI, CORBA, DCOM, some home-brewed protocol built on top of sockets and streams, or anything else. The client simply doesn't need to care about network protocols, because it can talk to the well-known interface that the service object implements. The service object takes care of any necessary communication on the network.

The client talks to the service through a well-known interface

Fig 3.Interface Implementation Different implementations of the same service interface can use completely different implementation approaches and completely different network protocols. A service can use specialized hardware to fulfill client requests, or it can do all its work in software. In fact, the implementation approach taken by a single service can evolve over time. The client can be sure it has a service object that understands the current implementation of the service, because the client receives the service object (by way of the lookup service) from the service provider itself. To the client, a service looks like the well-known interface, regardless of how the service is implemented.

4. Performance Analysis:

16 Sun Microsystems provides a development platform for E-business suit over 5,000 out of 13,000+ Companies worldwide, including Sun Corporate IT.Jini is one of the market leader among Solaris OE Servers,Storage,Java,J2EE,JME.Some of the facts and achievements associated with Jini are mentioned below :

Company 4D Networks, Inc. eko systems, Inc.

Industry

Product/ Category

Abstract

Telecom/Service Provider

Healthcare/Insurance

Java and Technologies

Operating Systems, Java and Technologies

Based on Jini network technology, 4D Pocket System uses the Jini lookup service to register services to its directory. Jini technology provided the system, called Frontiers, with the ability to connect to any type of medical equipment-including physiologic monitors, ventilators, and anesthesia machines without requiring clinicians to configure data or software drivers. Jini, a new paradigm for the development and management of distributed systems, provides mechanisms that enable systems to plug together to form an impromptu community. This practice report demonstrates how Jini can be applied in an industrial environment, or more exactly, how it can be used in the integration of embedded devices on-board trains in the backoffice IT infrastructure of railway operators. The results are encouraging and prove that Jini is the appropriate technology to link application servers and service gateways in embedded servers. This paper describes the architecture of a Jini based global computing environment, whose aim is to support the execution of sequential and parallel programs on a large scale. We describe the architecture, supported programming models of the system, as well as our current, ongoing work and future development plans.

Txomin Nieva, Swiss Federal Institute of Technology and ABB Corporate Research Ltd. Andreas Fabri, ABB Corporate Research Ltd. Abdenbi Benammour , Swiss Federal Institute of Technology

IEEE

Jini Technology Applied to Railway Systems

Zoltan Juhasz, Arpad Andics, Szabolcs Pota

2nd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID'02) May 21 - 24, 2002 Berlin, Germany

JM: A Jini Framework for Global Computing

6. Conclusion

As we have seen in this report, Jini attempts to raise the level of abstraction for distributed systems programming, from the network protocol level to the object interface level. In the emerging proliferation of embedded devices connected to networks, many pieces of a distributed system may come from different vendors. Jini makes it unnecessary for vendors of devices to agree on network level protocols that allow their devices to interact. Instead, vendors must agree on Java interfaces through which their devices can interact. The processes of discovery, join, and lookup, provided by the Jini runtime infrastructure, will enable devices to locate each other on the network. Once they locate each other, devices will be able to communicate with each other through Java interfaces.

7. References:

<http://www.jini.org> <http://www.sun.com/jini/licensing>. <http://wwws.sun.com/software/jini/specs/>
<http://www.java.sun.com/developer/products/jini/index.jsp> <http://csdl.computer.org/comp/proceedings/doa/2000/0819/00/08190251abs.htm>

