White Paper

IP Telephony: The Five Nines Story

IP Telephony Availability Overview

Availability can be defined as the probability that a product or service will operate when needed. When you buy a network or a network service, you expect availability. A network that isn't available is more than annoying or frustrating, it's expensive. With increased dependency on the network for business critical applications and voice, a down network costs everyone valuable time and money.

As enterprise organizations migrate legacy voice services to IP telephony, a routine consideration is availability of the IP telephony solution. Everyone has high expectations for voice service availability. A common goal is to consistently achieve 99.999% availability with 5.3 minutes of downtime per year or less. Cisco Systems understands that high availability and reliability are absolute requirements for all mission critical applications, including IP telephony solutions. Through a combination of PDIO (planning, design, implementation and operations) best-practices, Cisco IP telephony solutions can support 99.999% uptime. "Five nines" (99.999 percent) availability is not only a desired goal, it is an achievable goal that is possible using the recommended high availability IP telephony solution discussed in this paper.

The purpose of this document is to detail how PBX availability is defined and to demonstrate that Cisco does have a high availability five nines solution. The perception is that legacy PBX solutions are five nines end-to-end.

The truth is that legacy PBX availability achieves five nines by defining availability in a way that limits impact from non-redundant components. Using a similar definition, availability of the Cisco IP telephony solution is carefully examined using industry accepted theoretical availability analysis. This analysis includes Telcordia (formerly Bellcore) MTBF analysis using the parts count method, software availability analysis, power availability analysis and network design availability analysis.

Using the definition for availability of a redundant core architecture supporting Cisco CallManagers and PSTN gateways theoretical network hardware availability was shown to be 99.99993% in the hardware reliability section of this paper. This was based on a CI infrastructure model with redundant Catalyst 6509 chassis, redundant power supplies and redundant supervisor modules for Cisco CallManager access switches. Theoretical software availability was computed based on estimated software forced reloads for early deployment software but was negligible due to redundant devices at each layer and fast estimated reload times for Cisco products (6 minutes). Power availability was 99.99962% and is based on theoretical power availability analysis done by APC Corporation, available at http://www.apcc.com/support. This level of power availability requires UPS systems with generator backup and UPS monitoring with four hour response UPS service and support. Theoretical network design availability is 99.99947% and is based on failover times between redundant components at each network layer including the Cisco CallManager and PSTN gateways. The analysis is provided in the design section later in this paper. Using normal serial availability analysis we compute overall availability as the product of these results.

(.9999993) * (.9999962) * (.9999947) = 99.999% availability

Clearly Cisco has a five nines solution for IP telephony in a campus LAN deployment model based on available theoretical availability analysis techniques. Actual availability may vary due to other factors that include link/carrier availability, network management processes and expertise. By reading the following definition of availability and process by which availability is theoretically determined, anyone should be convinced that this is possible. The reader should also be aware that Cisco is committed to providing high availability and improving availability for Cisco customers.

Legacy PBX Availability

What does 99.999% reliability mean to a customer? It means absolute minimum downtime in terms of existing calls dropped and inability to place calls. Five nines equates to less that 5.26 minutes of downtime in a year of operation. Clearly this is an admirable target, and in mission critical applications, such as voice networks, it may be a requirement. Legacy PBX vendors are sometimes able to achieve five nines of availability using a strict definition of outages that occur directly to a redundant PBX system. The definition does not include end-to-end calling where broken phones, wiring or lost power can be a factor. The definition, as you will see below, is primarily a factor of total system uptime of the PBX itself.

Reliability in PBX Designs

It is important to understand how five nines reliability is measured in the legacy PBX world. In order to do that, let’s take a look at the basic components of a PBX and how they relate to measuring five nines uptime. The two diagrams in Figure 1 below illustrate very simplistic views of nonredundant and redundant PBXs.

Figure 1
Redundant and Non-redundant PBXs

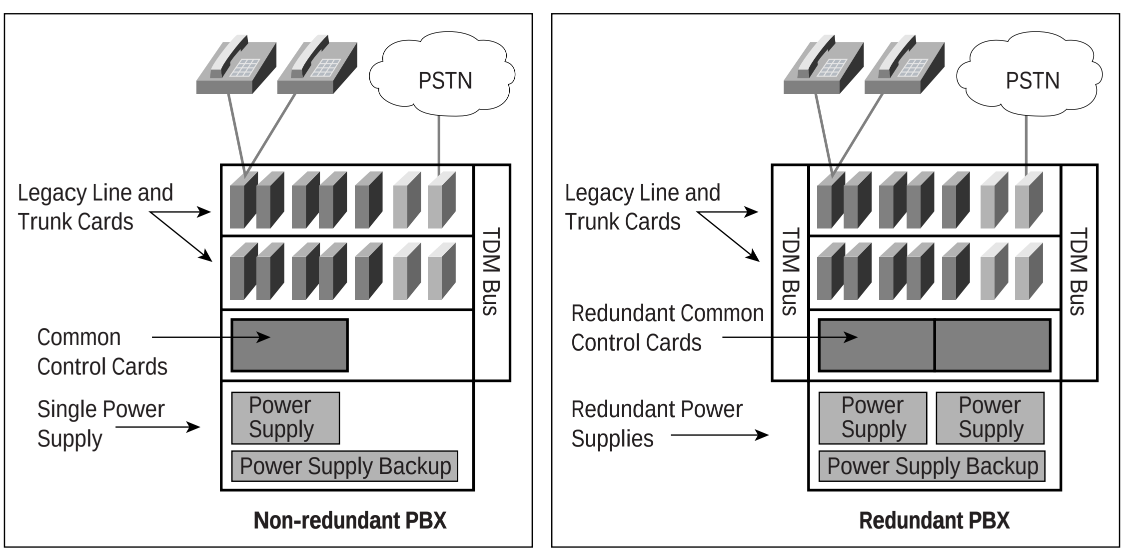

Description of Components:

- *Line Cards*—Analog line cards support analog devices (phones, modems, fax, recorders, etc.). Digital line cards support proprietary digital phones.
- *Trunk Cards*—Provide connectivity to the PSTN. Analog trunk cards for FXO, OPS, or TIE trunks. Digital trunk cards for T1, E1, PRI, BRI, etc.
- *Common Control*—CPU, memory, and other cards used to perform call processing, diagnostics, administration, etc., as well as common services (DTMF registers and senders for example).
- *TDM Bus*—This is the switching matrix for the PBX that connects all the cards together for control and voice traffic.
- *Power Supply*—Takes power from an external AC or DC source and distributes it to the shelves to power the cards.
- *Power Supply Backup*—This is an internal UPS that is found in some PBX's. It provides power to operate the system for several minutes, to provide protection from very brief power outages and to permit the system to go through an orderly shutdown. For more than just a few minutes of protection an external UPS system is required.

The main difference between non-redundant and redundant PBXs is optional duplication of the Common Control Cards, the Power Supplies, and the TDM Bus. Sometimes redundancy only means CPU and memory redundancy. Power supply redundancy is often a separately orderable redundancy option. Note, however, that the line and trunk cards are never redundant. No PBX vendor offers redundancy to the line card level.

Now that we've examined the components of a PBX, what types of failures would count in a calculation of five nines reliability? For example,

- Failure of a line card (typically 16 to 32 ports) putting phones out of service?
- Failure of a trunk card (typically 4 to 30 bearer channels) putting a group of trunks out of service?
- Loss of a couple of attendant consoles?

The answer? None of these would count as system downtime. In general, according to BellCore GR 512 standards, outages of 64 ports or less are not counted in the reliability calculations. The only thing that counts is total system downtime, which is loss of all lines, trunks, or all attendant consoles. In other words, in a non-redundant PBX you would have to lose a CPU or power supply; in redundant system you would have to lose both the primary and the backup element. That's why redundant PBXs are so reliable.

Of course, if you think about it logically, it would be difficult to consistently achieve 99.999% uptime if you counted minor failures, such as a non-redundant line card. Most PBX vendors offer four-hour response to minor outages such as loss of a line or trunk card (unless special service level agreements are in place). For instance if a PBX has 48 line cards in a PBX and four hour replacement is possible, the MTBF would have to be about 400,000 hours or 45 years. This uses the basic calculation of availability equaling MTBF / MTBF + MTTR. 400,000 hours is achievable, but is probably not implemented in most current PBX installations.

In fact, legacy PBX vendors are currently hard pressed to show five nines reliability for nonredundant systems. Again, five nines reliability typically infers that PBXs have maximum redundancy and power protection built in.

The legacy voice service provider has additional concerns including the local wiring system and phone reliability where redundancy, power backup and fast repair times are not possible. When end-to-end availability is closely examined the theoretical availability is less than five nines. In fact, many voice service providers target residential service availability in the range of 99.95%. This is closer to four hours of downtime per year and could be attributed to local power, broken phones or broken connections to the main switch. The main PBX may very well be five nines according to the above definition but individual users will experience lower availability due to issues outside of the switch itself.

In summary, legacy PBX systems can maintain five nines of availability. However, this level of availability only includes the PBX itself and does not reflect the true end-to-end voice availability because it does not include non-redundant line cards failures, power failures, broken phones and wiring systems. Many legacy voice service providers have budgeted end-to-end availability more in the range of 99.95% availability or 4+ hours of downtime per year.

IP Telephony Availability Definition

In a manner similar to legacy PBX vendors, Cisco can achieve five nines of availability with the IP Telephony solution by defining availability in terms of a redundant core system. Using the legacy PBX availability definition, Cisco can define Campus IP telephony availability as the ability to maintain and originate calls within the redundant core and distribution layers using a redundant Cisco CallManager cluster.

Using this definition, any component or individual software failure will force a failover to an alternate path or Cisco CallManager device. With one level of redundancy and four hour MTTR (mean time to repair), the IP telephony core can achieve five nines of availability in a manner analogous to the legacy PBX definition. This is shown with more detail in the following section.

Notice that Cisco has removed the access layer from the definition. Access switches typically support around 120 users, fairly close to a line card on a legacy PBX system. In the legacy world, this type of outage has been removed from the overall availability definition and calculation. One would suppose the reason is that the legacy PBX vendor cannot achieve five nines of reliability by including the access layer. Like legacy PBX vendors, Cisco cannot achieve five nines without redundancy. Access switches typically have 60,000 to 150,000 hours MTBF (mean time between failure) as computed by the Telcordia (formerly Bellcore) "Parts Count Method."[1] Using this methodology we find that a typical access switch with a 100,000-hour MTBF and a 4-hour MTTR has theoretical availability of 99.996% or 21 minutes of downtime per year. Clearly, no vendor can easily supply five nines of availability end-to-end and in fact current legacy voice service providers target availability significantly lower, around 99.95%. The Cisco IP phone has less complicated hardware and does maintain MTBF over 400,000 hours with five nines plus reliability.

1. Cisco uses a 2X multiplier from the Telcordia computed value to provide theoretical results in supports full line with historical reliability data collected for Cisco equipment.

Figure 2

“Common Infrastructure” or AVVID

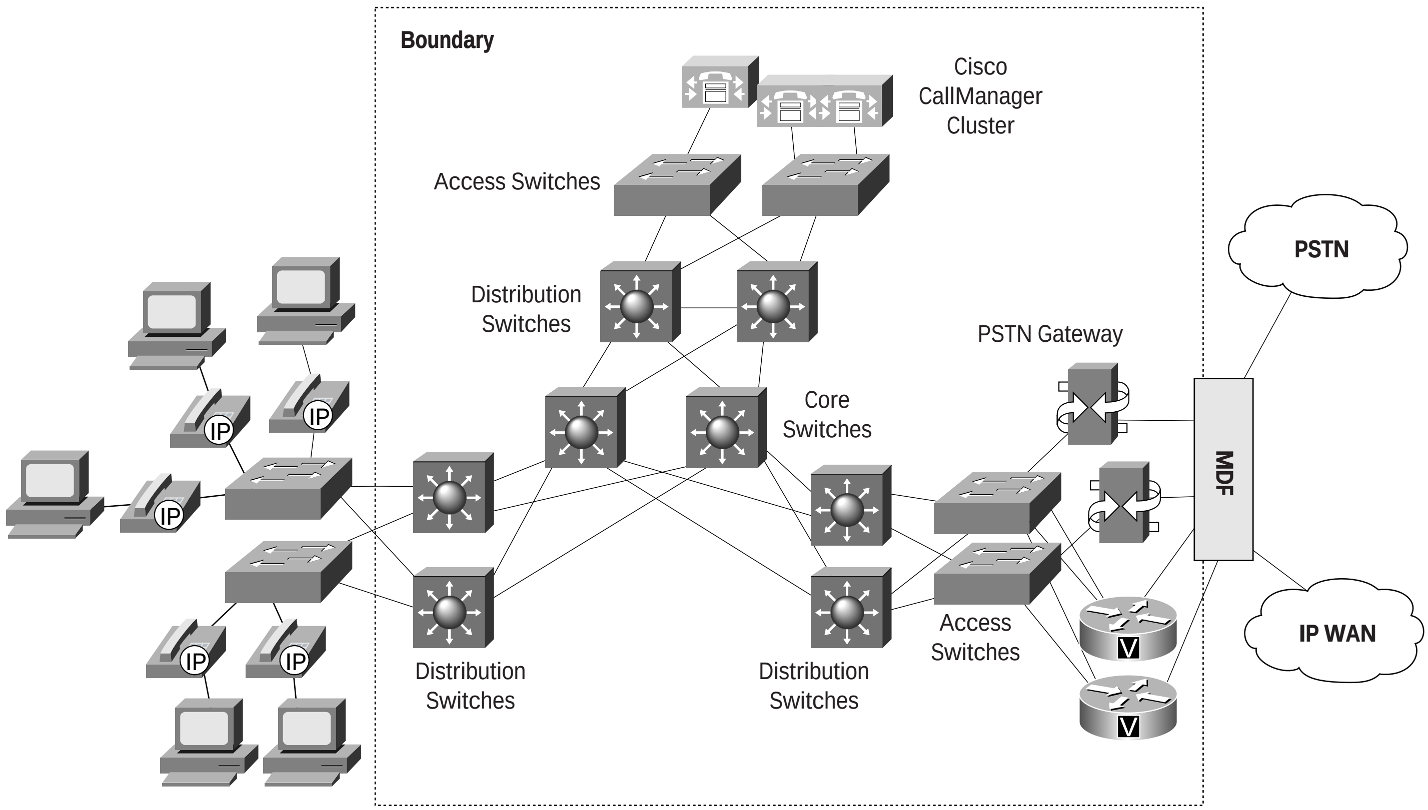

The diagram above shows a campus network core design often referred to as “Common Infrastructure” or AVVID (architecture for voice, video & integrated data). The design redundancy and can support five nines of availability using the Telcordia parts count method. More information on the Cisco five nines solution is available in the next section.

The IP Telephony Five Nines Solution

Availability of the IP telephony solution can be theoretically determined by closely examining availability characteristics in several areas and applying known availability analysis techniques.[2] By understanding and applying this methodology, organizations will better understand the availability characteristics of a potential solution. This is important to ensure the product or service provided meets expectations, to supply meaningful service level agreements, to create availability targets and to implement an availability improvement goal.

2. The book “High Availability Network Fundamentals” by Chris Oggerino provides detailed mathematical techniques for predicting availability of a network system.

A meaningful analysis of availability includes the following areas of network availability.

- Hardware reliability
- Software reliability
- Link/Carrier availability
- Power/Environment availability
- Network Design reliability
- User-error & network support processes

Link/Carrier availability may include all the areas as service provider networks are subject to the same failure root-causes as any enterprise environment. Network design includes topology and failover capabilities of the system. In IP telephony environments, this may include routing protocol convergence, Cisco CallManager system failover times or spanning tree convergence times. User-error & network support processes includes network management practices that promote network consistency, supportability and rapid problem resolution times.

This section provides a theoretical availability analysis of a campus IP telephony environment where N+1 redundancy is implemented for the core, distribution and Cisco CallManager server devices. Access devices are not included in the theoretical analysis. Link/Carrier connectivity network support processes are not evaluated for this LAN model. Figure 3 below shows the load sharing redundant topology of the five nines solution with devices outside the dotted line removed from the availability analysis.

Figure 3

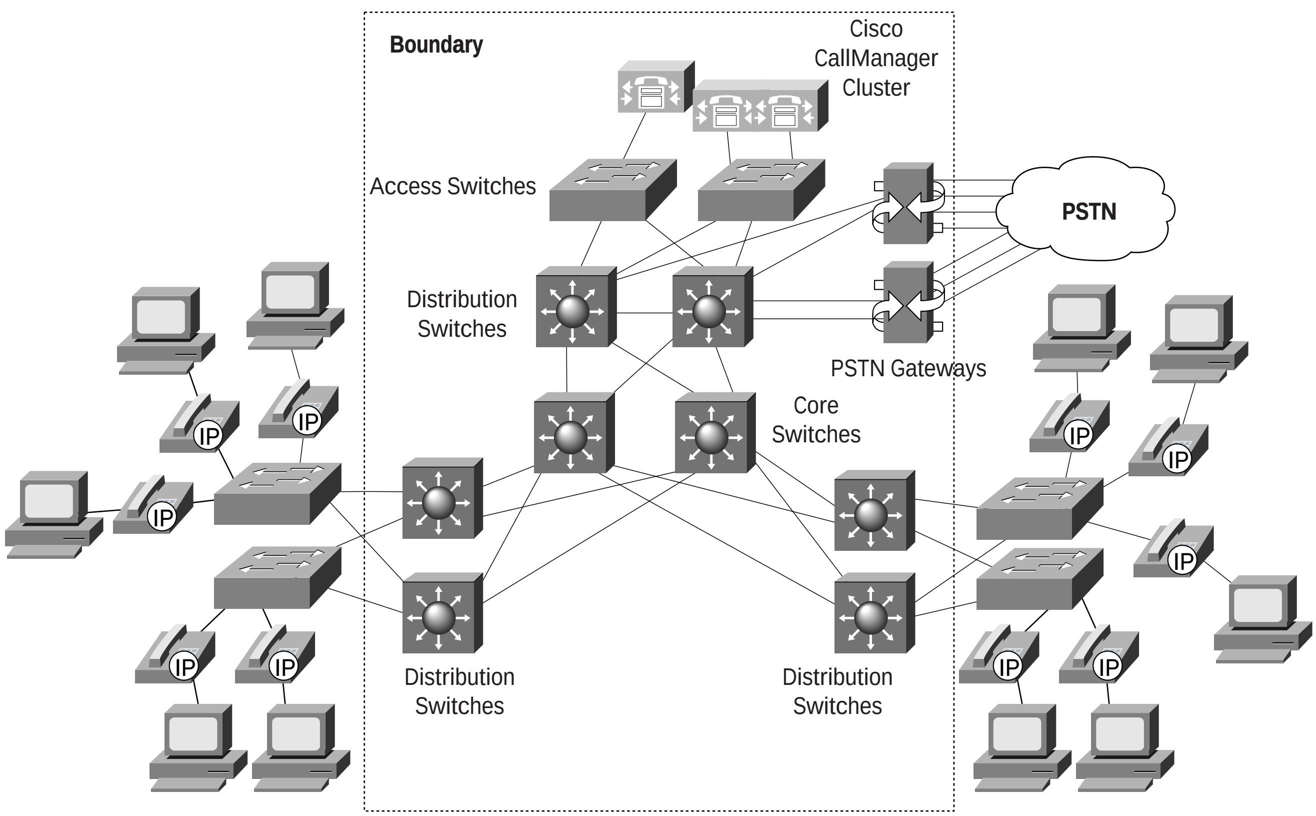

Hardware Reliability

Hardware reliability is estimated using the Telcordia "Parts Count Method" and applying the Cisco 2X factor based on historical failure data. This is accomplished by understanding MTBF values for devices, device components and paths through the network. A theoretical MTTR of four hours is used for the calculation and is considered to be the norm for theoretical vendor-comparison availability prediction analysis.

Using this methodology, hardware reliability was calculated to be 99.99993% for the IP telephony solution using the availability definition in the above section. The following information provides supporting detail for the hardware reliability analysis. Also keep in mind that many vendors based availability solely on hardware reliability.

Hardware analysis is accomplished by analyzing MTBF information along a path from point A to point B. In IP telephony environments, two paths are required for operation. One from phone to Cisco CallManager and one from phone to phone or phone to PSTN gateway. To allow for the two paths, analysis of two paths will be done individually and then together to determine hardware reliability in a typical CI recommended infrastructure model. Devices or device pairs that are on both paths need only to be calculated once. Note that the second path below includes gateways. This is not required for phone-to-phone availability within the local environment.

Figure 4
Reliability Blocks for Path Analysis

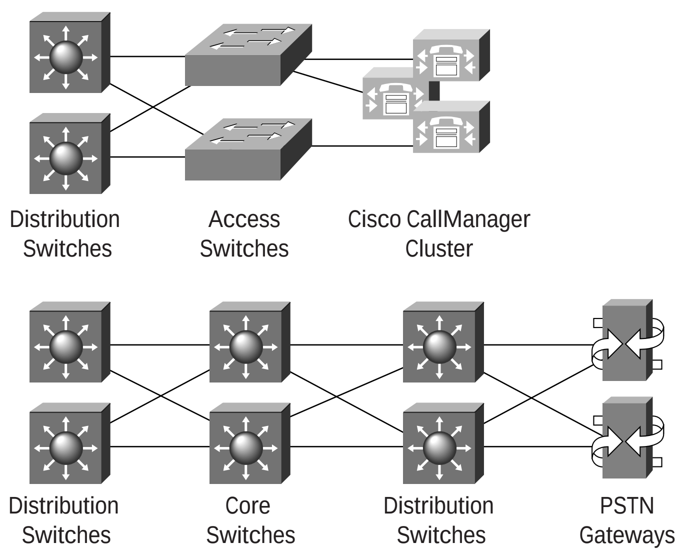

The reliability calculation begins with the calculation for one switch. The following tables show MTBF information for the modules and chassis of one distribution switch. MTBF and availability of the switch is identified via availability calculations identified in the book "High Availability Networking Fundamentals". The following equation is used to determine overall availability and can be explained in English as "for the components 'I' going from 1 to n (the number of components), multiply by availability i. In addition, the second table provides failure rate information. This is the theoretical percentage of time where the switch will fail during a one-year interval.

$$\text{System Availability} = \prod_{i=1}^{n} = \text{availability}_{(i)}$$

Table 1 Distribution Switch Module Reliability

Part #	Quantity	Quantity Required	MTBF	MTTR	Availability
WS-6509	1	1	369,897	4hr	99.99892%
WS-CAC-1300W	2	1	316,456	4hr	99.99999%
WS-X6K-SUP1A-MSFC2	1	1	46,235	4hr	99.99135%
WS-X6408A-BGIC	1	1	93,457	4hr	99.99572%

Table 2 Distribution Switch Reliability (Non-Redundant)

Availability	**99.985988%**
Annual Downtime	**73.7 minutes**
MTBF	**28,545 hrs**
Annual Failure Rate	**26.44%**

The next calculation can be applied to show availability of a redundant parallel system. This method only accounts for the second failure while the repair from the first failure is going on. Since the first repair only takes four hours, the availability looks very high. In reality, this almost never happens. The real problems are first, that even when there is a successful protocol switchover, all existing calls will be dropped and no new calls can be processed during the switchover. The availability due to this problem is addressed in the design section of this paper. Other problems include the inability to failover due to an unavailable or broken standby path or that the switchover does not occur because the failure is not detected. These are real problems that must be addressed via proper protocol design and testing, but for the most part have been shown to be successful with Cisco layer III routing protocol, HSRP and Spanning tree capabilities.

In this case of calculating basic hardware reliability we need to know the availability of redundant distribution switches. For this equation N equals the number of components in parallel (in this case 2) and I equals the component number. In our case, this is equivalent to 1 – [(1 – .99985988) * (1 – .99985988)] or 99.99999802%.

$$\text{Parallel availability} = 1 - \left[\prod_{i=1}^{n} (1 - \text{component availability}_{(i)}) \right]$$

Table 3 Distribution Switch Reliability (Redundant)

Availability	**99.99999804%**
Annual Downtime	**< 6 seconds**
MTBF	**101,880,673**
Annual Failure Rate	**.01%**

Using this methodology we can determine reliability of each redundant block and then use a normal series methodology to determine overall hardware reliability of the complete IP telephony path. In our case we have redundant blocks for two distribution layers, one core layer, Cisco CallManager layer and gateway layer. Overall availability is then determined by factoring each layer. But first we will need to determine module availability, chassis availability, and redundant chassis availability for each network layer. The following tables provide this information using the standard methods used above. Following this, availability of the complete path is provided.

Table 4 Core Switch Module Reliability

Part #	Quantity	Quantity Required	MTBF	MTTR	Availability
WS-6509	1	1	369,897	4hr	99.99892%
WS-CAC-1300W	2	1	316,456	4hr	99.99999%
WS-X6K-SUP1A-MSFC2	1	1	46,235	4hr	99.99135%
WS-X6408A-BGIC	2	2	93,457	4hr	99.99572%

Table 5 Core Switch Reliability (Non-redundant)

Availability	**99.981709%**
Annual Downtime	**96.2 minutes**
MTBF	**21,866 hrs**
Annual Failure Rate	**33.03%**

Table 6 Core Switch Reliability (Redundant)

Availability	**99.99999665%**
Annual Downtime	**< 6 seconds**
MTBF	**59,787,111**
Annual Failure Rate	**.01%**

Table 7 Access Switch Module Reliability

Part #	Quantity	Quantity Required	MTBF	MTTR	Availability
WS-6509	1	1	369,897	4hr	99.99892%
WS-CAC-1300W	2	1	316,456	4hr	99.99999%
WS-X6K-SUP1A-2GE	2	1	82,654	4hr	99.99999%
WS-G5484	2	1	432,470	4hr	99.99999%
WS-X6248A (10/100 card)	1	1	94,684	4hr	99.99577%

Table 8 Access Switch Reliability (Non-redundant)

Availability	**99.99469402%**
Annual Downtime	**27.9 minutes**
MTBF	**75,380 hrs.**
Annual Failure Rate	**10.98%**

Table 9 Access Switch Reliability (Redundant)

Availability	**99.99999972%**
Annual Downtime	**< 6 seconds**
MTBF	**710,343,430**
Annual Failure Rate	**.01%**

Table 10 Gateway/Access Switch Module Reliability

Part #	Quantity	Quantity Required	MTBF	MTTR	Availability
WS-6509	1	1	369,897	4hr	99.99892%
WS-CAC-1300W	2	1	316,456	4hr	99.99999%
WS-X6K-SUP1A-2GE	2	1	82,654	4hr	99.99999%
WS-G5484	2	1	432,470	4hr	99.99999%
WS-X6608-T1	1	1	40,512	4hr	99.99577%

Table 11 Gateway/Access Switch Reliability (Non-redundant)

Availability	**99.989045%**
Annual Downtime	**57.6 minutes**
MTBF	**36,511**
Annual Failure Rate	**21.34%**

Table 12 Gateway/Access Switch Reliability (Redundant)

Availability	**99.99999880%**
Annual Downtime	**< 6 seconds**
MTBF	**166,668,151**
Annual Failure Rate	**.01%**

The only remaining theoretical availability numbers pertain to the Cisco CallManager servers. Hardware is currently manufactured by outside vendors and the Telcordia MTBF information is not available. Cisco can only base availability for this product on the actual return rate of the product. The MCS 7835 currently has over 7 million hours of combined operation with a MTBF of over 1,000,000 hrs based on the number of returns during a 3-month period. Given the excellent current reliability of the MCS-7835 chassis and modules, the Cisco High Availability Services group offers the following estimates of availability for primary and redundant Cisco CallManager systems. The second table below provides reliability for redundant Cisco CallManager devices, which is a capability with the IP telephony solution.

Table 13 Cisco CallManager Chassis Reliability (Non-redundant)

Availability	**99.99%**
Annual Downtime	**57.6 minutes**
MTBF	**40,000**
Annual Failure Rate	**19.68%**

Table 14 Cisco CallManager Chassis Reliability (Redundant)

Availability	**99.99999%**
Annual Downtime	**< 6 seconds**
MTBF	**200,040,000**
Annual Failure Rate	**.01%**

Information is now available for the full path analysis. The following information is factored into the total path analysis calculation, which is obtained by multiplication of availability numbers for each redundant module. Given high availability at each layer overall hardware reliability of the IP telephony solution in a campus environment with an MTTR of four hours is 99.99993%. Again, this is determined simply by using the basic serial calculation method of each of the seven redundant modules described above.

Table 15 Redundant Layer Availability

Availability Path Analysis Based on Redundant (N+1) Modules						
Distribution 1—IP phone	Core	Distribution 2—PSTN access	PSTN access gateways	Distribution 3—Cisco CallManager	Access layer—Cisco CallManager	Cisco CallManager platform
99.99999	**99.99999**	**99.99999**	**99.99999**	**99.99999**	**99.99999**	**99.99999**

Figure 5

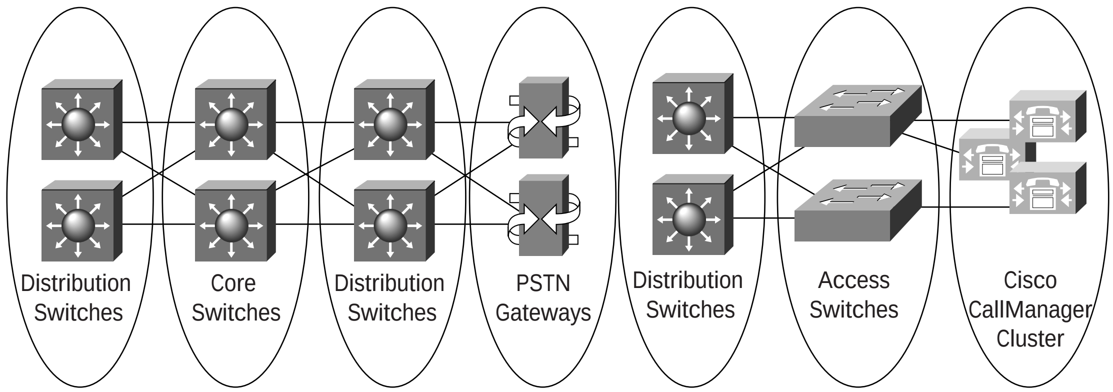

Software Reliability

To accurately predict the availability of a network or device, the software contribution to downtime must be considered. The problem however is that no industry standard exists for software reliability. Several methods are available for calculating software reliability but these methods are known to be inaccurate and difficult to perform. The best method known is to analyze field measurements for software reliability. This section examines a few Cisco internal field measurements and predicts software reliability based on techniques documented in the book "High Availability Networking Fundamentals" by Chris Oggerino.

Chris Oggerino and Scott Cherf, both of Cisco Systems, have been actively collecting and analyzing software reliability for over two years. Both individuals examined the number of software forced crashes from large sample sets with a variety of Cisco IOS versions. Oggerino and Cherf also standardized on a six minute MTTR as a conservative estimate for the time it would take a router to reboot and reestablish neighbors and traffic flows in a medium sized enterprise environment. These measurements did not account for software defects that did not restart the device, as these were difficult to measure. Using their methodology, Oggerino and Cherf conclude that Cisco IOS General Deployment (GD) software is five nines reliable. Cherf, for instance, concludes that Cisco IOS Software GD Release 11.1 is 99.99986% reliable with a MTBF of 71,740 hours and that the non-GD version of Release 11.1 is 99.9996% reliable with an MTBF of 25,484 hours. Both individuals also showed that software maturity was a major factor in software reliability.

Every Cisco customer should already know that software reliability is primarily a factor of software maturity. As the software ages and gains acceptance, defects are reduced until the software is considered stable with no problems expected. As IP telephony software increases in maturity, Cisco fully expects general deployment highly reliable software. Cisco is also committed to releasing five nines software by conducting extensive release tests on new software. Unfortunately, some problems surface with customer traffic patterns in real world environments. Cisco therefore recommends that all early deployment software be tested and/or piloted in the real-world environment to mitigate risk. Best-practices for change management and new solution deployment are detailed in the Cisco white papers at http://www.cisco.com/warp/public/126/index.shtml.

Because the Cisco IP telephony solutions utilize Cisco IOS software classified as "early deployment," a conservative estimate is needed that reflects this status. Based on field measurements for new software, a conservative estimate is created for this analysis of 10,000 hrs. MTBF and six minute MTTR. Oggerino uses this number in his book as an "extremely conservative number."[3] A conservative estimate is also created for General deployment software of

30,000 hours. These expected MTBF values will actually impact availability more due to failover. In the network design section, a 10,000 hour MTBF is used for newer and more complicated layers including gateway, Cisco CallManager and distribution layer. A 30,000-hour MTBF is used for more reliable software in the core and access layer. This analysis is provided in the design section of the paper later on.

We can use the software reliability number in the hardware calculation by factoring in software reliability as a component or module of each device. This can be done because software forced crashes or software restarts are independent activities isolated to one hardware platform. Field measurements also confirm that software failures do not occur to both redundant devices in a particular network layer at the same time. This is not to say that software defects cannot affect both devices in the same network layer. Routing protocols, HSRP, spanning tree and other protocols can affect availability of a network layer and should be designed carefully and tested prior to deployment. Fortunately, these protocols have been around for years and are not expected to cause non-availability under reasonable design conditions.

Using the previous definition of availability and fully redundant system, software does not impact availability of the system to a significant extent. For instance, hardware reliability of a distribution switch is 99.98598%. Software reliability using a 10,000 hour MTBF and six minute MTTR is 99.999%. This changes reliability of an individual switch to 99.98498%. Reliability of a redundant layer does not change significantly but drops less than a second per year from 99.99999804% to 99.99999774%. Since all calculations have been rounded off to five decimal places, this does not appear to impact availability. In the design section, we will use this MTBF however to look at convergence times of devices and the number of expected failovers to understand how this impacts availability.

Power/Environmental Availability

Power and environment are also potential impacts to availability that affect the overall availability of the IP telephony solution. Power is also unique in that it does not impact one device at a time like software or hardware. Its affects, or can affect, an entire building or multiple buildings at a time. This can impact all devices in the IP telephony availability definition including distribution, core, gateway and Cisco CallManager all at once. The calculations therefore change from device based to entire network based, which create a significant impact to theoretical availability depending on the power protection strategy used for IP telephony.

To understand theoretical availability for power and the environment, Cisco has turned to APC Corporation for theoretical availability.[4] APC can be contacted, http://www.apcc.com, for more precise availability information in North America based on geographical location, weather patterns, site details and installation specifics. In general however, the following characteristics apply to North American power. The data is variable from site to site and some areas like Florida are expected to experience an order of magnitude higher of non-availability. The data is also expected to be representative of Japan and Western Europe.

- The average number of outages sufficient to cause IT system malfunction per year at a typical site is approximately 15
- 90% of the outages are less than five minutes in duration
- 99% of the outages are less than one hour in duration
- Total cumulative outage duration is approximately 100 minutes per year

3. Page 67 of “High Availability Networking Fundamentals” by Chris Oggerino

4. Based on technical note #24 from APC Corporation located at http://sturgeon.apcc.com/kbasewb2.nsf/Tnotes+External/6F45EFB214C9DDFF8525672300568C B6?OpenDocument

APC also provides theoretical availability based on the power protection strategy and concludes that availability levels of five nines or higher require a UPS system with a minimum of one hour batterylife or a generator with a onsite service contract or four hour response for UPS system failures or problems. The Cisco recommended high availability solution requires additional support to achieve five nines overall. The HA IP telephony solution therefore must include UPS and generator backup for all distribution, core, gateway and Cisco CallManager devices. In addition, the organization should have UPS systems that have auto-restart capability and a service contract for four hour response to support the UPS/generator. Given the following recommended core infrastructure, APC corporation estimates that non-availability will be two minutes per year.

IP telephony high availability power/environment recommendations:

- UPS & generator backup
- UPS systems have auto-restart capability
- UPS system monitoring
- Four hour service response contract for UPS system problems
- Maintain recommended equipment operating temperatures 24X7

Overall power availability using the above solution is estimated to be 99.99962%. This impacts overall availability, 99.99993%, in the same way a new module would affect a device in a serial system. The calculation used to determine overall estimated availability is then (.9999962) * (.9999993) or 99.99955%.

Network Design Availability

Network design availability includes network topology and network protocol design used to provide scalability, performance and convergence following a device or link failure. This is critical to availability. Without a properly designed network with recommended protocols and configuration, convergence or failover times could easily exceed five minutes, eliminating any opportunity to achieve a five nines solution due to expected software and hardware failures. To achieve the highest availability in network design, protocols need to determine when a failure has occurred and immediately switch to a redundant component or device. The system must also recognize and monitor backup components to ensure that the failover will be successful. Fortunately, Cisco is the recognized leader in these networking capabilities and can provide high availability as long as the design and implementation meet IP telephony design rules. This section helps to identify the design rules for the defined IP telephony availability solution and also estimates availability based on failure rates seen in software and hardware.

The following design rules apply to the Cisco IP telephony five nines solution:

- Standard Common Infrastructure design with core/distribution/access layers
- EIGRP or OSPF routing in the campus with default timers
- HSRP for all IP access subnets including phones and Cisco CallManager servers
- Spanning tree environments eliminated or fast-converging spanning tree features (Use spanning tree only when multiple access switches are needed on 1 subnet/VLAN to support HA services)
- Dual home access switches to two distribution switches
- Dual home distribution switches to two core switches
- In general do not trunk devices at the same layer (except where multiple access switches are needed on 1 subnet/VLAN to support HA devices)
- Quality of Service configured across the LAN (or plenty of bandwidth)

When the recommended Common Infrastructure design for Campus networks is implemented for IP telephony we can investigate the failure rate at each network layer. This is accomplished by first understanding the time required for failover and convergence at each network layer. Fortunately, this has been previously tested within Cisco. The next step is to understand the annual failure rate a component at each layer for both software and hardware. The failure rate can then be multiplied by the expected failover time to understand the overall impact to availability. A major assumption in this analysis is that the network has adequate resources (bandwidth, CPU & processor memory) to handle the entire load while the alternate device is recovering. Figure 6 below can be used to identify the layers where failover will occur.

Figure 6
IP Telephony Path

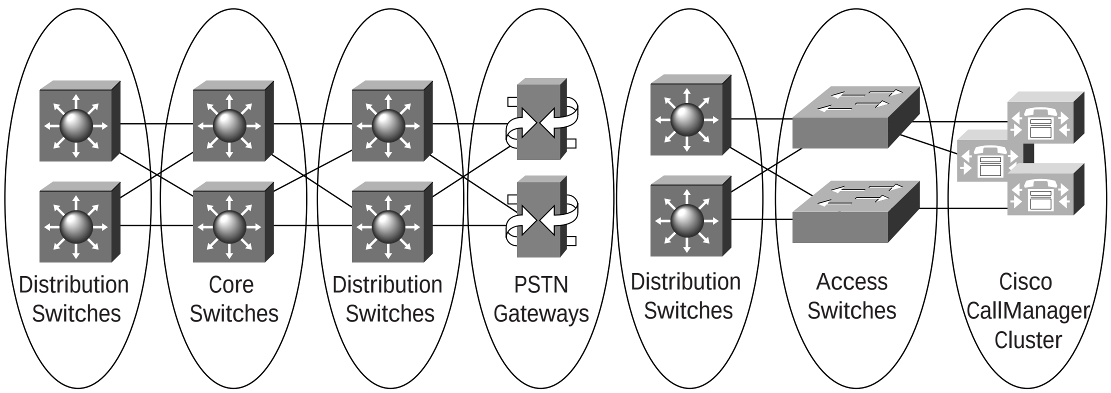

Of the seven layers above, four separate failover mechanisms will be discussed that lead to slightly different failover or convergence times. The following table provides detail on failover and convergence requirements and how they occur.

Table 16 Failover Analysis

Failover Type	Analysis
Distribution layer	A failed distribution device or links may include three different failover mechanisms. These include HSRP, IP routing and possibly spanning tree. Thom Bryant documents testing results for failover in the “Redundant Network Design Failure Analysis” paper available internally at Cisco. Thom has applied tuning to improve network convergence at the distribution layer. For the purpose of this document, we are simply taking a conservative number based on his test results, 30 seconds.
Core layer	Distribution devices should have equal cost routes through the core, making IP convergence unnecessary to route traffic. Thom Bryant documents testing results for failover in the “Redundant Network Design Failure Analysis” white paper. Thom has applied tuning to improve network convergence at the distribution layer. For the purpose of this document, we are simply taking a conservative number based on his test results, 10 seconds. The document is currently available.
PSTN gateway	The Cisco CallManager will attempt to initiate a PSTN call through the designated route group, device list or alternate device list. If the Cisco CallManager cannot initiate a call, the alternate device is selected. This takes approximately 10 seconds.

Table 16 Failover Analysis

Failover Type	Analysis
Cisco CallManager and access switches	Redundant Cisco CallManagers are not currently stateful. This capability is expected in future releases. Currently, a TCP keepalive is sent every 30 seconds. Three simultaneous negative responses cause failover to the standby, which can initiate new calls. Current calls are not dropped during failover however DTMF signaling info is lost and call may be reset. A failed access device will cause a failover to the redundant Cisco CallManager. Because three keepalives can occur in 61 seconds, the mean time before the 1st keepalive is used bringing the average downtime to 75 seconds. Shortly following this, the redundant Cisco CallManager will have the ability to accept new call requests.

Table 17 Expected Failure Rate Per Year

	Distribution	Core	Access	Gateway	Cisco CallManager
Software MTBF	10,000 hrs	30,0000 hrs	30,000 hrs.	10,000 hrs 1	0,000 hrs
Hardware MTBF (from Table 2)	28,545 hrs	21,866 hrs	75,380 hrs	36,511 hrs	40,000 hrs
Software failure rate per year	.876 (365*24/ 10,000 = .876)	.292	.292	.876	.876
Hardware failure rate per year	.306	.401	.116	.240	.219
Total failure rate[1]	1.182	.693	.408	1.116	1.095
Failure impact	30 seconds	10 seconds	75 seconds	10 seconds	75 seconds
Estimated down time per year	35 seconds	7 seconds	31 seconds	11 seconds	82 seconds
Total downtime	166 seconds	2.77 minutes			
Network design availability	99.99947%				

1. It is generally not a good idea to combine failure rates since software and hardware failure rates have different MTTR values. However this only applies to switchover time due to a failed component.

Availability has now been created for hardware, software, power/environment & network design. Keep in mind that the theoretical calculation has ignored the number of dropped calls. This is currently difficult to estimate. Calls in fact are not necessarily dropped but can be perceived as such when failover takes over a few seconds.

Overall power availability was estimated to be 99.99962% and hardware/software availability was expected to be .9999993. Using the estimate for network design availability of 99.99947 we can use a serial model to calculate theoretical availability. This is (.9999962) * (.9999947) * (.9999993) which equate to 99.999% or approximately 5.3 minutes of downtime per year.

User Error and Network Support Processes

User error and network support process can account for significant downtime in any IT infrastructure.[5] However, the amount of downtime that an organization will incur due to user error or lack of processes is difficult to theoretically determine. The best-known method is to analyze expertise, staffing and processes to determine gaps that may impact availability. Unfortunately, this is well beyond the scope of this document. For this purpose of theoretically estimating availability, user-error and process have not been included. This does not mean that it should be ignored. Both traditional PBX vendors and IP telephony vendors experience similar user-error and process issues and therefore should be addressed in both environments.

Rather than focus on the impact to availability in this area, it makes sense to at least identify some of the user-error and process issues that may impact availability in an IP telephony environment. These processes should promote the ability to deploy solutions successfully, quickly identify and resolve faults, keep the network modular & consistent, and provide the ability to resolve potential problems proactively. More information is available on these processes in the IP telephony solutions guide at: http://www.cisco.com/warp/customer/788/solution_guide/

5. A Gartner Group report states that 40% of non-availability is due to user-error and process.

New technology deployment planning & design

Staffing & expertise

Operational support plan (Service repair plans & problem priorities, contacts, escalations)

Fault monitoring & detection

Cisco CallManager & device backups

Performance management/Capacity planning

Cisco CallManager and device configuration backups

Software & application management (versions control)

Change Management

Lab validation

New technology pilot

Phone provisioning

MAC (move, add, change) process

IP/DHCP management

Telephone inventory management

Dial plan management

Fraud detection

Disaster recovery

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: 65 317 7777
Fax: 65 317 7799

201640/ETMG 02/02