Hepting et al v. AT&T Corp. et al Doc. 277 Att. 4

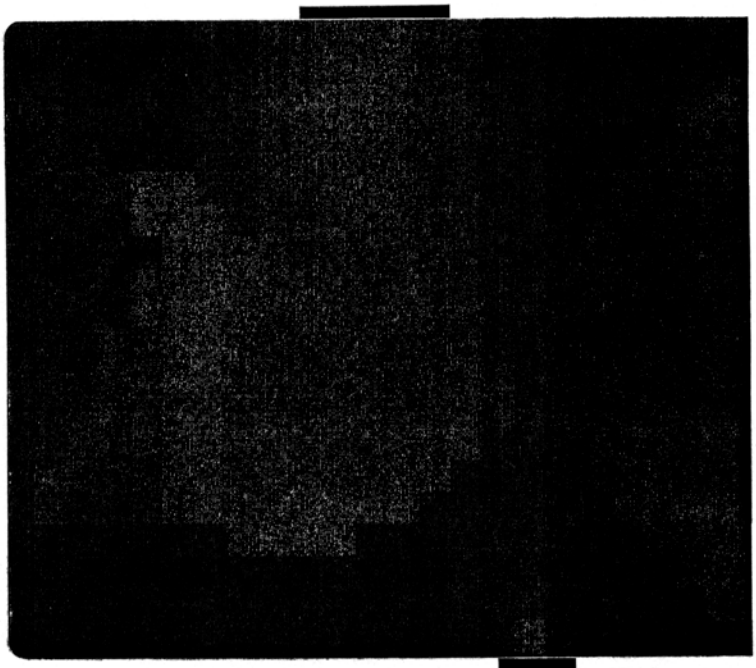

85. In addition to its real time capabilities, the [redacted] offering can subsequently analyze large volumes of data in order to reconstruct session content as needed from the captured collections of packets. This would include e-mail, web browsing, voice over IP (VoIP), and other common kinds of Internet communication.[36]

86. It would, in my judgment, be an error to evaluate the capabilities of this configuration – substantial though they are – solely on the basis of the equipment deployed by AT&T to the [redacted] Room. The AT&T documents clearly indicate the presence of an [redacted] network, apparently operating at [redacted].[37] This network, while much smaller than AT&T's CBB Internet backbone network, is nonetheless quite substantial.

87. The [redacted] backbone was logically distinct from the AT&T Common Backbone (CBB), but this does not necessarily mean that it had dedicated physical transmission facilities. It most probably operated over AT&T's standard optical fiber-based transmission systems, but using different high speed services – in effect, different circuits – than the CBB. If this network were carrying nothing more than a subset of AT&T's normal commercial traffic, they might not have

[36] [redacted]

[37] Klein Exh. C, pp. 6, 12, 42.

Dockets.Justia.com

felt the need to do more – it has long been considered permissible to transmit *Sensitive but Unclassified Information (SUCI)* over separate fiber-based transmission paths. Had there been greater sensitivity about the data, it might have been protected in other ways, for instance by means of link encryption.

88. The obvious and natural design for a massive surveillance system for IP-based data, and the one most cost-effective to implement, would in my judgment be comprised of the following elements: (1) massive data capture at the locations where the data can be tapped, (2) high speed screening and reduction[38] of the captured data at the point of capture in order to identify data of interest, (3) shipment of the data of interest to one or two central collection points for more detailed analysis, and (4) intensive analysis and cross correlation of the data of interest by very powerful processing engines at the central location or locations. The AT&T documents demonstrate that equipment that is well suited for the first three of these tasks was deployed to [redacted] and, with high probability, to other locations. I infer that the fourth element also exists at one or more locations.

89. Staff to analyze the data would probably be based at the central locations. There would be no need to station analysts (as distinct from field support personnel) in the [redacted] rooms where the data was collected. It is likely that the data were directly available for analysis by staff of the agency that funded the [redacted] (which runs counter to normal practice in the case of CALEA); otherwise, there would have been no need for a private [redacted], separate from the CBB.

90. The [redacted] technology could potentially be used in a number of different ways, some of which could be welfare-enhancing. The concern that must be raised in this case is that, in conjunction with the diversion of large volumes of traffic described in the Klein Declaration and the Klein Exhibits, this configuration appears to have the capability to enable surveillance and analysis of Internet content on a massive scale, including both overseas and purely domestic traffic.

[38] The [redacted] appears to be ideally suited to this role. It is, as previously noted, designed to apply a large collection of tests against a huge volume of data at very high speed.

TRAFFIC CAPTURED AT [redacted] ROOM

91. In this section, I explain my conclusions about the volume and type of communications traffic gathered by the [redacted] Room [redacted].

92. The Klein Declaration and Klein Exhibits B & C describe traffic diversions associated with fiber-based circuits in the [redacted] facility.

93. All of the diverted data pertains to AT&T's Common Backbone (CBB), the IP-based network that supports AT&T's Internet access customers, and that also carries AT&T's VoIP services (voice over the Internet).[39] Nothing in the documents suggests that conventional telephony traffic was diverted to the [redacted] Configuration.

94. The last page of Klein Exhibit B provides a list of CBB [redacted] (defined below) links that were to be split and diverted to the [redacted] Configuration.

95. Nothing in the documents suggests that AT&T's *on net* traffic – traffic from one AT&T customer to another – was diverted at the time. AT&T may at some point in time have made some provision for its international customers (whose traffic to other AT&T customers would also be on net), but the documents provide no guidance. My assumption is that on net traffic was not diverted during the time frame to which the documents pertain.

96. Before proceeding, it is helpful to introduce and clarify some terms. *Peering* is the process whereby Internet providers interchange traffic destined for their respective customers, and for customers of their customers. The Network Reliability and Interoperability Council (NRIC), an advisory panel to the FCC, defined peering in this way:[40]

> *Peering* is an agreement between ISPs to carry traffic for each other and for their respective customers. Peering does not include the obligation to carry traffic to third

---

[39] See *In the Matter of AT&T Petition for Declaratory Ruling that AT&T's Phone-to-Phone IP Telephony Services are Exempt from Access Charges,* FCC WC Docket 02-361, Petition of AT&T, at 24 (filed Oct. 18, 2002), at http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6513386921 (Exhibit Q).

[40] Report of the NRIC V Interoperability Focus Group, an advisory panel to the FCC: "Service Provider Interconnection for Internet Protocol Best Effort Service," page 7, available at http://www.nric.org/fg/fg4/ISP_Interconnection.doc (Exhibit R). *See also* chapter 14 of Marcus, *Designing Wide Area Networks and Internetworks: A Practical Guide*, Addison Wesley, 1999 (Exhibit S).

> parties. Peering is usually a bilateral business and technical arrangement, where two providers agree to accept traffic from one another, and from one another's customers (and thus from their customers' customers) . . . .

97. In the figure below, AT&T and Backbone B are *peers*. They have agreed to exchange traffic for their respective customers. Traffic from AT&T customer 1 to AT&T customer 2 is *on net* traffic – it remains on AT&T's network. Traffic from AT&T customer 1 to customer 3 (a customer of backbone B) is *off net* traffic.

FIGURE 4

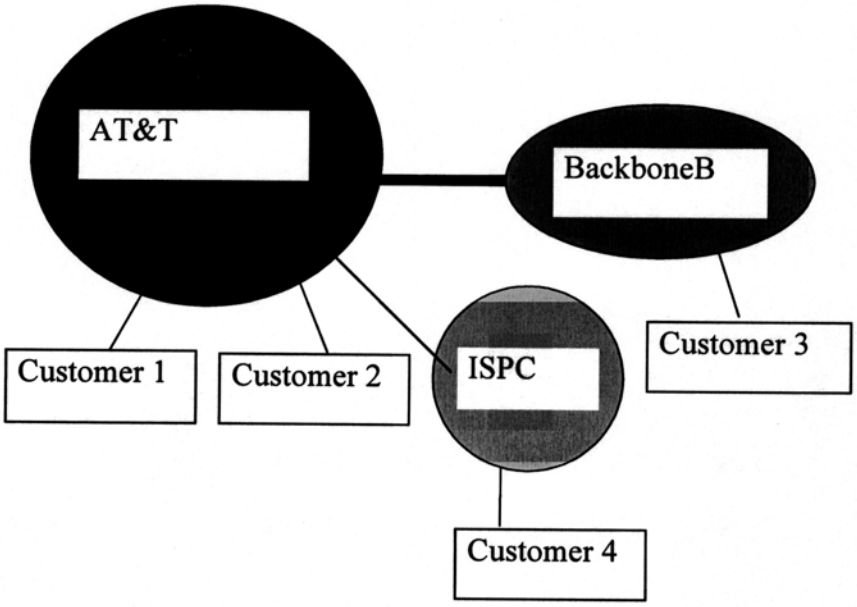

98. In the figure, ISP C is a *transit customer* of AT&T. ISP C pays AT&T to carry its traffic, not only to AT&T customers, but to customers of other ISPs as well (such as, for example, Customer 3). In the context of this discussion, AT&T can regard traffic from Customer 4 to Customers 1 and 2 as being on net, in the sense that it does not traverse a peering connection.

99. It is perhaps also worth noting that AT&T and its peers and their many transit customers do not merely connect to the Internet; rather they *are* the Internet. The Internet is not a single, huge and over-arching network, but rather a collection of independent networks that collectively comprise a worldwide communications stratum.

100. Again, the last page of Exhibit B provides a list of CBB [redacted] that were to be split and diverted to the [redacted] Configuration. The sizes of these circuits are listed, with some at [redacted], some at [redacted], and some at [redacted]. These

are all quite substantial circuits – the [redacted] are apparently on a par with the largest circuits that were in widespread use in AT&T's CBB Internet backbone at the time.

101. Traffic to and from several very large Internet providers at that time [redacted] was delivered over OC-48 circuits. Traffic to and from another group of large providers [redacted] was delivered over [redacted] circuits. Traffic to and from smaller, but still quite substantial, providers [redacted] was delivered over [redacted] circuits.

102. Large Internet backbone providers typically use direct interconnects (private peering) to exchange traffic with their largest "trading partners in bits," the firms with which they exchange the largest volume of traffic. For providers where the volume of traffic exchange at some location is large enough to warrant peering arrangements, but not large enough to justify the cost of a separate circuit for private peering, it is customary instead to interconnect with multiple peers at a so-called "public peering point" in order to exchange traffic with multiple providers there.[41] AT&T was connected to [redacted] configuration.

103. At the point where I left Genuity in July 2001 (some eighteen months before these splitters were deployed), I was intimately familiar with our traffic exchange patterns with other providers. Our measurement instrumentation ranked with the very best in the industry at that time. It is possible to draw many inferences about traffic flows among other providers from one's own traffic exchanges.

104. Based on my experience at Genuity, I believe that the traffic that was diverted represented all, or substantially all, of AT&T's peering traffic in the [redacted]

105. I base my reasoning on the knowledge of Genuity's peering traffic patterns, and on my general understanding of peering traffic patterns in the industry. As of July 2001, our three largest peers were WorldCom, AT&T and Sprint, collectively representing 50-60% of our traffic.

---

[41] *See* Marcus, *Designing Wide Area Networks and Internetworks: A Practical Guide*, Addison Wesley, 1999, pages 280-282 (Exhibit S).